Containment for Rule-Based Ontology-Mediated Queries

Pablo Barceló
Center for Semantic Web Research &
DCC, University of Chile
pbarcelo@dcc.uchile.cl

Gerald Berger
Institute of Information Systems
TU Wien
gberger@dbai.tuwien.ac.at

Andreas Pieris
School of Informatics
University of Edinburgh
apieris@inf.ed.ac.uk

arXiv:1703.07994v3 [cs.DB] 19 Apr 2017

ABSTRACT

Many efforts have been dedicated to identifying restrictions on ontologies expressed as tuple-generating dependencies (tgds), a.k.a. existential rules, that lead to the decidability for the problem of answering ontology-mediated queries (OMQs). This has given rise to three families of formalisms: guarded, non-recursive, and sticky sets of tgds. In this work, we study the containment problem for OMQs expressed in such formalisms, which is a key ingredient for solving static analysis tasks associated with them. Our main contribution is the development of specially tailored techniques for OMQ containment under the classes of tgds stated above. This enables us to obtain sharp complexity bounds for the problems at hand, which in turn allow us to delimitate its practical applicability. We also apply our techniques to pinpoint the complexity of problems associated with two emerging applications of OMQ containment: distribution over components and UCQ rewritability of OMQs.

1. INTRODUCTION

Motivation and goals. The novel application of knowledge representation tools for handling incomplete and heterogeneous data is giving rise to a new field, recently coined as *knowledge-enriched data management* [6]. A crucial problem in this field is *ontology-based data access* (OBDA) [51], which refers to the utilization of ontologies (i.e., sets of logical sentences) for providing a unified conceptual view of various data sources. Users can then pose their queries solely in the schema provided by the ontology, abstracting away from the specifics of the individual sources. In OBDA, one interprets the ontology Σ and the user query q, which is typically a *union of conjunctive queries* (UCQ), or, equivalently, the expressions defined by the select-project-join-union operators of relational algebra, as two components of one composite query $Q = (\mathbf{S}, \Sigma, q)$, known as *ontology-mediated query* (OMQ); $\mathbf{S}$ is called the *data schema*, indicating that Q will be posed on databases over $\mathbf{S}$ [19]. Therefore, OBDA is often realized as the problem of answering OMQs.

Following recent work [24, 26, 27, 40], we focus on the case where the ontology is defined by a set of *tuple-generating dependencies* (tgds), a.k.a. *existential rules* or *Datalog$^{\pm}$ rules*. Handling such OMQs implies new challenges for classical database tasks. Interestingly, some of these challenges are by now well-studied; most notably (a) *query evaluation* [8, 24, 25, 27]: given an OMQ $Q = (\mathbf{S}, \Sigma, q)$, a database D over $\mathbf{S}$, and a tuple of constants $\bar{c}$, does $\bar{c}$ belong to the evaluation of q over every extension of D that satisfies Σ, or, equivalently, is $\bar{c}$ a *certain answer* for Q over D? and (b) *relative expressiveness* [19, 42, 43]: how does the expressiveness of OMQs compare to the one of other query languages? Surprisingly, despite its prominence, no work to date has carried out an in-depth investigation of *containment* for OMQs based on tgds and UCQs.

Query containment is a fundamental static analysis task that amounts to check if the evaluation of a query is always contained in the evaluation of another query. Several database tasks crucially depend on the ability to check query containment; these include, e.g., query optimization, view-based query answering, querying incomplete databases, integrity checking, and implication of dependencies: cf. [22, 30, 36, 37, 39, 45]. A particularly important instance of the containment problem is the one defined by the class of CQs. It follows from the seminal work of Chandra and Merlin [29] that CQ containment is polynomially equivalent to CQ evaluation, and thus NP-complete. The NP upper bound is not affected if we consider UCQs [54]. This is seen as a positive result for practical applications that rely on UCQ containment, as the input (the two UCQs) is small. In addition, it shows a stark difference with more expressive relational query languages, e.g., relational algebra (or, equivalently, first-order logic), for which containment is undecidable.

The main goal of this work is to understand up to which extend the good computational properties of UCQ containment discussed above can be leveraged to the containment problem for OMQs based on tgds and UCQs (simply called OMQs from now on). In particular, we want to understand which classes of tgds guarantee the decidability of the problem, and, whenever this is the case, how can we obtain complexity bounds that are reasonable for practical purposes. We also want to understand what is the exact relationship between OMQ containment and evaluation for such classes. Let us stress that, apart from the traditional applications of containment mentioned above, it has been recently shown that OMQ containment has applications on other important static analysis tasks for OMQs, namely, distribution over components [15], and UCQ rewritability [16].

The context. As one might expect, when considered in its full generality, i.e., without any restrictions on the set of tgds, the OMQ containment problem is undecidable. To understand, on the other hand, which restrictions lead to decidability, we recall the two main reasons that render the general containment problem undecidable. These are:

Undecidability of query evaluation: OMQ evaluation is, in general, undecidable [12], and it can be reduced to OMQ containment. More precisely, OMQ containment is undecidable whenever query evaluation for at least one of the involved languages (i.e., the language of the left-hand or the right-hand side query) is undecidable.

Undecidability of containment for Datalog: decidability of query evaluation does not ensure decidability of query containment. A prime example is Datalog, or, equivalently, the OMQ language based on *full* tgds. Datalog containment is undecidable [55], and thus, OMQ containment is undecidable if the involved languages extend Datalog.

In view of the above observations, we focus on languages that (a) have a decidable query evaluation, and (b) do not extend Datalog. The main classes of tgds, which give rise to OMQ languages with the desirable properties, can be classified into three main families depending on the underlying syntactic restrictions: (i) *guarded* tgds [24], which contain inclusion dependencies and linear tgds, (ii) *non-recursive* sets of tgds [35], and (iii) *sticky* sets of tgds [27].

While the decidability of containment for the above OMQ languages can be established via translations into query languages with a decidable containment problem, such translations do not lead to optimal complexity upper bounds (details are given below). Therefore, the main goal of our paper is to develop specially tailored decision procedures for the containment problem under the OMQ languages in question, and ideally obtain precise complexity bounds. Our second goal is to exploit such techniques in the study of distribution over components and UCQ rewritability of OMQs.

Our contributions. The complexity of OMQ containment for the languages in question is given in Table 1. Using small fonts, we recall the complexity of OMQ evaluation in order to stress that containment is, in general, harder than evaluation. We divide our contributions as follows:

Linear, non-recursive and sticky sets of tgds. The OMQ languages based on linear, non-recursive, and sticky sets of tgds share a useful property: they are *UCQ rewritable* (implicit in [40]), that is, an OMQ can be rewritten into a UCQ. This property immediately yields decidability for their associated containment problems, since UCQ containment is decidable [54]. However, the obtained complexity bounds are not optimal, since the UCQ rewritings are unavoidably very large [40]. To obtain more precise bounds, we reduce containment to query evaluation, an idea that is often applied in query containment; see, e.g., [29, 31, 54].

Consider a UCQ rewritable OMQ language $\mathbb{O}$. If Q_1 and Q_2 belong to $\mathbb{O}$, both with data schema $\mathbf{S}$, then we can establish a *small witness property*, which states that non-containment of Q_1 in Q_2 can be witnessed via a database over $\mathbf{S}$ whose size is bounded by an integer $k \geq 0$, the maximal size of a disjunct in a UCQ rewriting of Q_1. For linear tgds, such an integer k is polynomial, but for non-recursive and sticky sets of tgds it is exponential (implicit in [40]). The above small witness property allows us to devise a simple non-deterministic algorithm, which makes use of query evaluation as a subroutine for checking non-containment of Q_1 in Q_2: guess a database D over $\mathbf{S}$ of size at most k, and then check if there is a certain answer for Q_1 over D that is not a certain answer for Q_2 over D. This algorithm allows us to obtain optimal upper bounds for OMQs based on linear and sticky sets of tgds; however, the exact complexity of OMQs based on non-recursive sets of tgds remains open:

- For OMQs based on linear tgds, the problem is in PSpace, and in Π_2^P if the arity is fixed. The PSpace-hardness is shown by reduction from query evaluation [47], while the Π_2^P-hardness is inherited from [17].
- For OMQs based on sticky sets of tgds, the problem is in coNExpTime, and in Π_2^P if the arity of the schema is fixed. The coNExpTime-hardness is shown by exploiting the standard tiling problem for the exponential grid, while the Π_2^P-hardness is inherited from [17].
- Finally, for OMQs based on non-recursive sets of tgds, containment is in ExpSpace and hard for $\mathrm{P}^{\mathrm{NEXP}}$, even for fixed arity. The lower bound is shown by exploiting a recently introduced tiling problem [34].

We conclude that in all these cases OMQ containment is harder than evaluation, with one exception: the OMQs based on linear tgds over schemas of unbounded arity.

Guarded tgds. The OMQ language based on guarded tgds is not UCQ rewritable, which forces us to develop different tools to study its containment problem. Let us remark that guarded OMQs can be rewritten as guarded Datalog queries (by exploiting the translations devised in [9, 43]), for which containment is decidable in 2ExpTime [20]. But, again, the known rewritings are very large [43], and hence the reduction of containment for guarded OMQs to containment for guarded Datalog does not yield optimal upper bounds.

To obtain optimal bounds for the problem in question, we exploit *two-way alternating parity automata on trees* (2WAPA) [32]. We first show that if Q_1 and Q_2 are guarded OMQs such that Q_1 is not contained in Q_2, then this is witnessed over a class of "tree-like" databases that can be represented as the set of trees accepted by a 2WAPA $\mathfrak{A}$. We then build a 2WAPA $\mathfrak{B}$ with exponentially many states that recognizes those trees accepted by $\mathfrak{A}$ that represent witnesses to non-containment of Q_1 in Q_2. Hence, Q_1 is contained in Q_2 iff $\mathfrak{B}$ accepts no tree. Since the emptiness problem for 2WAPA is feasible in exponential time in the number of states [32], we obtain that containment for guarded OMQs is in 2ExpTime. A matching lower bound, even for fixed arity schemas, follows from [16].

Similar ideas based on 2WAPA have been recently used to show that containment for OMQs based on expressive *description logics* (DLs) is in 2ExpTime [16]. In the DL context, schemas consist only of unary and binary relations. Our automata construction, however, is different from the one in [16] for two reasons: (a) we need to deal with higher arity relations, and (b) even for unary and binary relations, our OMQ language allows to express properties that are not expressible by the DL-based OMQ languages studied in [16].

Combining languages. The above complexity results refer to the containment problem relative to a certain OMQ language $\mathbb{O}$, i.e., both queries fall in $\mathbb{O}$. However, it is natural

	Arbitrary Arity	Bounded Arity
Linear	PSPACE-c PSPACE-c	Π_2^P-c NP-c
Sticky	coNEXPTIME-c EXPTIME-c	Π_2^P-c NP-c
Non-recursive	in EXPSPACE and P^{NEXP}-hard NEXPTIME-c	in EXPSPACE and P^{NEXP}-hard NEXPTIME-c
Guarded	2EXPTIME-c 2EXPTIME-c	2EXPTIME-c EXPTIME-c

Table 1: Complexity of OMQ containment – in small fonts, we recall the complexity of OMQ evaluation.

to consider the version of the problem where the involved OMQs fall in different languages. Unsurprisingly, if the left-hand side query is expressed in a UCQ rewritable OMQ language (based on linear, non-recursive or sticky sets of tgds), we can use the algorithm that relies on the small witness property discussed above, which provides optimal upper bounds for almost all the considered cases (the only exception is the containment of sticky in non-recursive OMQs over schemas of unbounded arity). Things are more interesting if the ontology of the left-hand side query is expressed using guarded tgds, while the ontology of the right-hand side query is not guarded. By exploiting automata techniques, we show that containment of guarded in non-recursive OMQs is in 3EXPTIME, while containment of guarded in sticky OMQs is in 2EXPTIME. We establish matching lower bounds, even over schemas of fixed arity, by refining techniques from [31].

Applications. Our techniques and results on containment for guarded OMQs can be applied to other important static analysis tasks, in particular, distribution over components and UCQ rewritability.

The notion of distribution over components has been introduced in [3], in the context of declarative networking, and it states that the answer to an OMQ Q can be computed by parallelizing it over the (maximally connected) components of the database. If this is the case, then Q can always be evaluated in a distributed and coordination-free manner. The problem of deciding distribution over components for OMQs has been recently studied in [15]. However, the exact complexity of the problem for guarded OMQs has been left open. By exploiting our results on containment, we can show that it is 2EXPTIME-complete.

It is well-known that the OMQ language based on guarded tgds is not UCQ rewritable. In view of this fact, it is important to study when a given guarded OMQ Q can be rewritten as a UCQ. This has been studied for OMQs based on central Horn DLs [16, 18]. Interestingly, our automata-based techniques for guarded OMQ containment can be adapted to decide in 2EXPTIME whether an OMQ based on guarded tgds over unary and binary relations is UCQ rewritable; a matching lower bound is inherited from [16]. Our result generalizes the result that deciding UCQ rewritability for OMQs based on $\mathcal{ELHI}$, one of the most expressive members of the $\mathcal{EL}$-family of DLs, is 2EXPTIME-complete [16].

Discussion on Applicability. As shown in Table 1, the containment problem for OMQs based on linear sets of tgds is PSPACE-complete, and thus can be solved in single-exponential time. This is not a big practical drawback since the containment problem corresponds to a static analysis task. In fact, the runtime is single exponential only in the size of the UCQs and the maximum arity of the underlying schema, which are typically very small. For such tasks, a single-exponential time procedure is considered to be acceptable, and it is actually the norm in many cases including database and verification problems; see, e.g., [1, 50, 52].

For OMQs based on sticky, non-recursive and guarded sets of tgds, the containment problem becomes coNEXPTIME-complete, P^{NEXP}-hard and 2EXPTIME-complete, respectively. This means that we require double-exponential time to solve the problem, which is practically not acceptable. Nevertheless, for sticky sets of tgds, the runtime is double-exponential only in the maximum arity of the schema, while for guarded sets of tgds is double-exponential only in the size of the UCQs and the maximum arity of the schema. This is good news since, as said above, the size of the UCQs and the arity are typically small, and usually UCQs in OMQs are much smaller than the ontologies.

For non-recursive sets of tgds, on the other hand, the runtime is double-exponential, not only in the maximum arity, but also in the number of predicates occurring in the ontology. It is unrealistic to assume that the number of predicates occurring in real-life ontologies is small. This fact, together with the fact that the precise complexity of OMQ containment for non-recursive sets of tgds is still open, suggests that a more careful complexity analysis is needed. This is left as an interesting open problem for future work.

Organization. Preliminaries are in Section 2. In Section 3 we introduce the OMQ containment problem. Containment for UCQ rewritable OMQs is studied in Section 4, and for guarded OMQs in Section 5. In Section 6 we consider the case where the involved queries fall in different languages. In Section 7 we discuss the applications of our results on guarded OMQ containment and we conclude in Section 8. Proofs and additional details can be found in the appendix.

2. PRELIMINARIES

Databases and conjunctive queries. Let $\mathbf{C}$, $\mathbf{N}$, and $\mathbf{V}$ be disjoint countably infinite sets of *constants*, *(labeled) nulls* and (regular) *variables* (used in queries and dependencies), respectively. A *schema* $\mathbf{S}$ is a finite set of relation symbols (or predicates) with associated arity. We write R/n to denote that R has arity n. A *term* is a either a constant, null or variable. An *atom* over $\mathbf{S}$ is an expression of the form $R(\bar{v})$, where $R \in \mathbf{S}$ is of arity $n > 0$ and $\bar{v}$ is an n-tuple of terms. A *fact* is an atom whose arguments consist only of constants. An *instance* over $\mathbf{S}$ is a (possibly infinite) set of atoms over $\mathbf{S}$ that contain constants and nulls, while a *database* over $\mathbf{S}$ is a finite set of facts over $\mathbf{S}$. We may call an instance and a database over $\mathbf{S}$ an $\mathbf{S}$-instance and $\mathbf{S}$-database, respectively.

The *active domain* of an instance I, denoted $dom(I)$, is the set of all terms occurring in I.

A *conjunctive query* (CQ) over $\mathbf{S}$ is a formula of the form:

$$q(\bar{x}) \ := \ \exists \bar{y}\big(R_1(\bar{v}_1) \wedge \cdots \wedge R_m(\bar{v}_m)\big), \tag{1}$$

where each $R_i(\bar{v}_i)$ $(1 \leq i \leq m)$ is an atom without nulls over $\mathbf{S}$, each variable mentioned in the $\bar{v}_i$'s appears either in $\bar{x}$ or $\bar{y}$, and $\bar{x}$ are the free variables of q. If $\bar{x}$ is empty, then q is a *Boolean CQ*. As usual, the evaluation of CQs is defined in terms of homomorphisms. Let I be an instance and $q(\bar{x})$ a CQ of the form (1). A *homomorphism* from q to I is a mapping h, which is the identity on $\mathbf{C}$, from the variables that appear in q to the set of constants and nulls $\mathbf{C} \cup \mathbf{N}$ such that $R_i(h(\bar{v}_i)) \in I$, for each $1 \leq i \leq m$. The *evaluation of $q(\bar{x})$ over I*, denoted $q(I)$, is the set of all tuples $h(\bar{x})$ of constants such that h is a homomorphism from q to I. We denote by $\mathbb{CQ}$ the class of conjunctive queries.

A *union of conjunctive queries* (UCQ) over $\mathbf{S}$ is a formula of the form $q(\bar{x}) := q_1(\bar{x}) \vee \cdots \vee q_n(\bar{x})$, where each $q_i(\bar{x})$ is a CQ of the form (1). The *evaluation of $q(\bar{x})$ over I*, denoted $q(I)$, is the set of tuples $\bigcup_{1 \leq i \leq n} q_i(I)$. We denote by $\mathbb{UCQ}$ the class of union of conjunctive queries.

Tgds and the chase procedure. A *tuple-generating dependency* (tgd) is a first-order sentence of the form:

$$\forall \bar{x} \forall \bar{y}\big(\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z}\, \psi(\bar{x}, \bar{z})\big), \tag{2}$$

where ϕ and ψ are conjunctions of atoms without nulls. For brevity, we write this tgd as $\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z}\, \psi(\bar{x}, \bar{z})$ and use comma instead of $\wedge$ for conjoining atoms. Notice that ϕ can be empty, in which case the tgd is called *fact tgd* and is written as $\top \rightarrow \exists \bar{z}\, \psi(\bar{x}, \bar{z})$. We assume that each variable in $\bar{x}$ is mentioned in some atom of ψ. We call ϕ and ψ the *body* and *head* of the tgd, respectively. The tgd in (2) is logically equivalent to the expression $\forall \bar{x}(q_\phi(\bar{x}) \rightarrow q_\psi(\bar{x}))$, where $q_\phi(\bar{x})$ and $q_\psi(\bar{x})$ are the CQs $\exists \bar{y}\, \phi(\bar{x}, \bar{y})$ and $\exists \bar{z}\, \psi(\bar{x}, \bar{z})$, respectively. Thus, an instance I over $\mathbf{S}$ *satisfies* this tgd iff $q_\phi(I) \subseteq q_\psi(I)$. We say that an instance I satisfies a set Σ of tgds, denoted $I \models \Sigma$, if I satisfies every tgd in Σ. We denote by $\mathbb{TGD}$ the class of (finite) sets of tgds.

The *chase* is a useful algorithmic tool when reasoning with tgds [24, 35, 47, 49]. We start by defining a single chase step. Let I be an instance over a schema $\mathbf{S}$ and $\tau = \phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z}\, \psi(\bar{x}, \bar{z})$ a tgd over $\mathbf{S}$. We say that τ is *applicable* w.r.t. I if there exists a tuple $(\bar{a}, \bar{b})$ of terms in I such that $\phi(\bar{a}, \bar{b})$ holds in I. In this case, *the result of applying τ over I with $(\bar{a}, \bar{b})$* is the instance J that extends I with every atom in $\psi(\bar{a}, \bar{\perp})$, where $\bar{\perp}$ is the tuple obtained by simultaneously replacing each variable $z \in \bar{z}$ with a fresh distinct null not occurring in I. For such a single chase step we write $I \xrightarrow{\tau, (\bar{a}, \bar{b})} J$.

Let us assume now that I is an instance and Σ a finite set of tgds. A *chase sequence for I under Σ* is a sequence:

$$I_0 \ \xrightarrow{\tau_0, \bar{c}_0} \ I_1 \ \xrightarrow{\tau_1, \bar{c}_1} \ I_2 \ \cdots$$

of chase steps such that: (1) $I_0 = I$; (2) for each $i \geq 0$, τ_i is a tgd in Σ; and (3) $\bigcup_{i \geq 0} I_i \models \Sigma$. We call $\bigcup_{i \geq 0} I_i$ the *result* of this chase sequence, which always exists. Although the result of a chase sequence is not necessarily unique (up to isomorphism), each such result is equally useful for our purposes, since it can be homomorphically embedded into every other result. Thus, from now on, we denote by $chase(I, \Sigma)$ the result of an arbitrary chase sequence for I under Σ.

Ontology-mediated queries. An *ontology-mediated query* (OMQ) is a triple $(\mathbf{S}, \Sigma, q)$, where $\mathbf{S}$ is a schema, Σ is a set of tgds (the ontology), and q is a (U)CQ over $\mathbf{S} \cup sch(\Sigma)$ (and possibly other predicates), with $sch(\Sigma)$ the set of predicates occurring in Σ.[1] We call $\mathbf{S}$ the *data schema*. Notice that the set of tgds can introduce predicates not in $\mathbf{S}$, which allows us to enrich the schema of the UCQ q. Moreover, the tgds can modify the content of a predicate $R \in \mathbf{S}$, or, in other words, R can appear in the head of a tgd of Σ. We have explicitly included $\mathbf{S}$ in the specification of the OMQ to emphasize that it will be evaluated over $\mathbf{S}$-databases, even though Σ and q might use additional relational symbols.

The semantics of an OMQ is given in terms of certain answers. The *certain answers* to a UCQ $q(\bar{x})$ w.r.t. a database D and a set Σ of tgds is the set of tuples:

$$cert(q, D, \Sigma) \quad = \bigcap_{I \supseteq D, I \models \Sigma} \{\bar{c} \in dom(I)^{|\bar{x}|} \mid \bar{c} \in q(I)\}.$$

Consider an OMQ $Q = (\mathbf{S}, \Sigma, q)$. The *evaluation* of Q over an $\mathbf{S}$-database D, denoted $Q(D)$, is defined as $cert(q, D, \Sigma)$. It is well-known that $cert(q, D, \Sigma) = q(chase(D, \Sigma))$; see, e.g., [24]. Thus, $Q(D) = q(chase(D, \Sigma))$.

Ontology-mediated query languages. We write $(\mathbb{C}, \mathbb{Q})$ for the OMQ language that consists of all OMQs of the form $(\mathbf{S}, \Sigma, q)$, where Σ falls in the class $\mathbb{C}$ of tgds, i.e., $\mathbb{C} \subseteq \mathbb{TGD}$ (concrete classes of tgds are discussed below), and the query q falls in $\mathbb{Q} \in \{\mathbb{CQ}, \mathbb{UCQ}\}$. A problem that is quite important for our work is *OMQ evaluation*, defined as follows:

PROBLEM :	$\mathsf{Eval}(\mathbb{C}, \mathbb{Q})$
INPUT :	An OMQ $Q = (\mathbf{S}, \Sigma, q(\bar{x})) \in (\mathbb{C}, \mathbb{Q})$, an $\mathbf{S}$-database D, and $\bar{c} \in dom(D)^{\lvert\bar{x}\rvert}$.
QUESTION :	Does $\bar{c} \in Q(D)$?

It is well-known that $\mathsf{Eval}(\mathbb{TGD}, \mathbb{CQ})$ is undecidable; implicit in [12]. This has led to a flurry of activity for identifying syntactic restrictions on sets of tgds that make the latter problem decidable. Such a restriction defines a subclass $\mathbb{C}$ of tgds. The known decidable classes of tgds are classified into three main decidability paradigms, which, in turn, give rise to decidable OMQ languages:

Guardedness: A tgd is *guarded* if its body contains an atom, called *guard*, that contains all the body-variables. Although the chase under guarded tgds does not necessarily terminate, the problem of deciding whether a tuple of constants is a certain answer to a UCQ w.r.t. a database and a set of guarded tgds is decidable. This follows from the fact that the result of the chase has *bounded treewidth* (see, e.g., [24]). Let $\mathbb{G}$ be the class of (finite) sets of guarded tgds. Then:

PROPOSITION 1. [24] $\mathsf{Eval}(\mathbb{G}, \mathbb{CQ})$ *and* $\mathsf{Eval}(\mathbb{G}, \mathbb{UCQ})$ *are* 2EXPTIME*-complete, and* EXPTIME*-complete for fixed arity.*

An important subclass of guarded tgds is the class of *linear* tgds whose body consists of a single atom. We write $\mathbb{L}$ for the class of (finite) sets of linear tgds.

PROPOSITION 2. [25, 47] $\mathsf{Eval}(\mathbb{L}, \mathbb{CQ})$ *and* $\mathsf{Eval}(\mathbb{L}, \mathbb{UCQ})$ *are* PSPACE*-complete, and* NP*-complete for fixed arity.*

[1] OMQs can be defined for arbitrary first-order theories, not only tgds, and first-order queries, not only UCQs [19].

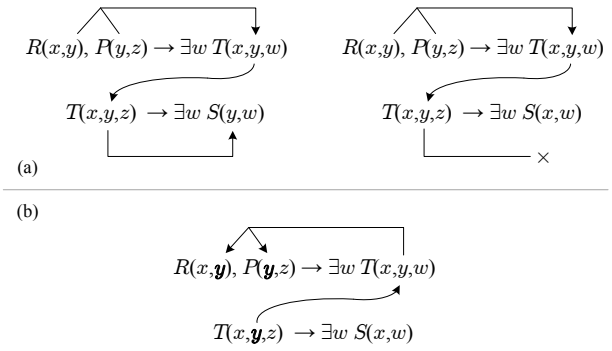

Figure 1: Stickiness and Marking.

Non-recursiveness: A set Σ of tgds is *non-recursive* (a.k.a. *acyclic* [35, 48]), if its predicate graph, the directed graph that encodes how the predicates of $sch(\Sigma)$ depend on each other, is acyclic. Non-recursiveness ensures the termination of the chase, and thus decidability of OMQ evaluation. Let $\mathbb{NR}$ be the class of non-recursive (finite) sets of tgds. Then:

PROPOSITION 3. [48] $\mathsf{Eval}(\mathbb{NR}, \mathbb{CQ})$ *and* $\mathsf{Eval}(\mathbb{NR}, \mathbb{UCQ})$ *are* NEXPTIME*-complete, even for fixed arity.*

Stickiness: This condition ensures neither termination nor bounded treewidth of the chase. Instead, the decidability of OMQ evaluation is obtained by exploiting query rewriting techniques (more details on query rewriting are given in Section 4). The goal of stickiness is to capture joins among variables that are not expressible via guarded tgds, but without forcing the chase to terminate. The key property underlying this condition can be described as follows: during the chase, terms that are associated (via a homomorphism) with variables that appear more than once in the body of a tgd (i.e., join variables) are always propagated (or "stick") to the inferred atoms. This is illustrated in Figure 1(a); the left set of tgds is sticky, while the right set is not. The formal definition is based on an inductive marking procedure that marks the variables that may violate the semantic property of the chase described above [27]. Roughly, during the base step of this procedure, a variable that appears in the body of a tgd τ but not in every head-atom of τ is marked. Then, the marking is inductively propagated from head to body as shown in Figure 1(b). Finally, a finite set of tgds Σ is *sticky* if no tgd in Σ contains two occurrences of a marked variable. Let $\mathbb{S}$ be the class of sticky (finite) sets of tgds. Then:

PROPOSITION 4. [27] $\mathsf{Eval}(\mathbb{S}, \mathbb{CQ})$ *and* $\mathsf{Eval}(\mathbb{S}, \mathbb{UCQ})$ *are* EXPTIME*-complete, and* NP*-complete for fixed arity.*

3. OMQ CONTAINMENT: THE BASICS

The goal of this work is to study in depth the problem of checking whether an OMQ Q_1 is *contained* in an OMQ Q_2, both over the same data schema $\mathbf{S}$, or, equivalently, whether $Q_1(D) \subseteq Q_2(D)$ over every (finite) $\mathbf{S}$-database D. In this case we write $Q_1 \subseteq Q_2$; we write $Q_1 \equiv Q_2$ if $Q_1 \subseteq Q_2$ and $Q_2 \subseteq Q_1$. The *OMQ containment* problem in question is defined as follows; $\mathbb{O}_1$ and $\mathbb{O}_2$ are OMQ languages $(\mathbb{C}, \mathbb{Q})$, where $\mathbb{C}$ is a class of tgds (e.g., linear, non-recursive, sticky, etc.), and $\mathbb{Q} \in \{\mathbb{CQ}, \mathbb{UCQ}\}$:

PROBLEM :	$\mathsf{Cont}(\mathbb{O}_1, \mathbb{O}_2)$
INPUT :	Two OMQs $Q_1 \in \mathbb{O}_1$ and $Q_2 \in \mathbb{O}_2$.
QUESTION :	Does $Q_1 \subseteq Q_2$?

Whenever $\mathbb{O}_1 = \mathbb{O}_2 = \mathbb{O}$, we refer to the containment problem by simply writing $\mathsf{Cont}(\mathbb{O})$.

In what follows, we establish some simple but fundamental results, which help to better understand the nature of our problem. We first investigate the relationship between evaluation and containment, which in turn allows us to obtain an initial boundary for the decidability of our problem, i.e., we can obtain a positive result only if the evaluation problem for the involved OMQ languages is decidable (e.g., those introduced in the previous section). We then focus on the OMQ languages introduced in Section 2 and observe that, once we fix the class of tgds, it does not make a difference whether we consider CQs or UCQs. In other words, we show that an OMQ in $(\mathbb{C}, \mathbb{UCQ})$, where $\mathbb{C} \in \{\mathbb{G}, \mathbb{L}, \mathbb{NR}, \mathbb{S}\}$, can be rewritten as an OMQ in $(\mathbb{C}, \mathbb{CQ})$. This fact simplifies our later complexity analysis since for establishing upper (resp., lower) bounds it suffices to focus on CQs (resp., UCQs).

3.1 Evaluation vs. Containment

As one might expect, OMQ evaluation and OMQ containment are strongly connected. In fact, as we explain below, the former can be easily reduced to the latter. But let us first introduce some auxiliary notation. Consider a database D and a tuple $\bar{c} = (c_1, \ldots, c_n) \in dom(D)^n$, where $n \geq 0$. We denote by $q_{D,\bar{c}}(\bar{x})$, where $\bar{x} = (x_{c_1}, \ldots, x_{c_n})$, the CQ obtained from the conjunction of atoms occurring in D after replacing each constant c with the variable x_c. Consider now an OMQ $Q = (\mathbf{S}, \Sigma, q(\bar{x})) \in (\mathbb{C}, \mathbb{CQ})$, where $\mathbb{C}$ is some class of tgds, an $\mathbf{S}$-database D, and a tuple $\bar{c} \in dom(D)^{|\bar{x}|}$. It is not difficult to show that

$$\bar{c} \in Q(D) \iff \underbrace{(sch(\Sigma), \varnothing, q_{D,\bar{c}})}_{Q_1} \subseteq \underbrace{(sch(\Sigma), \Sigma, q)}_{Q_2}.$$

Let $\mathbb{O}_\varnothing$ be the OMQ language that consists of all OMQs of the form $(\mathbf{S}, \varnothing, q)$, i.e., the set of tgds is empty, where q is a CQ. It is clear that $Q_1 \in \mathbb{O}_\varnothing$ and $Q_2 \in (\mathbb{C}, \mathbb{CQ})$. Therefore, for every OMQ language $\mathbb{O} = (\mathbb{C}, \mathbb{CQ})$, where $\mathbb{C}$ is a class of tgds, we immediately get that:

PROPOSITION 5. $\mathsf{Eval}(\mathbb{O})$ *can be reduced in polynomial time into* $\mathsf{Cont}(\mathbb{O}_\varnothing, \mathbb{O})$.

We now show that the problem of evaluation is reducible to the complement of containment. Let us say that, for technical reasons which will be made clear in a while, we focus our attention on classes $\mathbb{C}$ of tgds that are *closed under fact tgd extension*, i.e., for every set $\Sigma \in \mathbb{C}$, a set obtained from Σ by adding a (finite) set of fact tgds is still in $\mathbb{C}$. This is not an unnatural assumption since every reasonable class of tgds, such as the ones introduced above, enjoy this property. Consider now an OMQ $Q = (\mathbf{S}, \Sigma, q(\bar{x})) \in (\mathbb{C}, \mathbb{CQ})$, where $\mathbb{C}$ is some class of tgds, an $\mathbf{S}$-database D, and a tuple $\bar{c} \in dom(D)^{|\bar{x}|}$. It is easy to see that

$$\bar{c} \in Q(D) \iff \underbrace{(\mathbf{S}, \Sigma^\star_D, q^\star_{\bar{c}})}_{Q_1} \not\subseteq \underbrace{(\mathbf{S}, \varnothing, \exists x\, P(x))}_{Q_2},$$

where $\Sigma^\star_D$ is obtained from Σ by renaming each predicate R in Σ into $R^\star \notin \mathbf{S}$ and adding the set of fact tgds

$$\{\top \to R^\star(c_1, \ldots, c_k) \mid R(c_1, \ldots, c_k) \in D\},$$

$q^\star_{\bar{c}}$ is obtained from $q(\bar{c})$ by renaming each predicate R into $R^\star \notin \mathbf{S}$, and the predicate P does not occur in $\mathbf{S}$. Indeed, the above equivalence holds since $P \notin \mathbf{S}$ implies that $Q_2(D) = \varnothing$, for every $\mathbf{S}$-database D. Since $\mathbb{C}$ is closed under fact tgd extension, $Q_1 \in (\mathbb{C}, \mathbb{CQ})$, while $Q_2 \in \mathbb{O}_\varnothing$. We write $\mathsf{coCont}(\mathbb{O}_1, \mathbb{O}_2)$ for the complement of $\mathsf{Cont}(\mathbb{O}_1, \mathbb{O}_2)$. Hence, for every OMQ language $\mathbb{O} = (\mathbb{C}, \mathbb{CQ})$, where $\mathbb{C}$ is a class of tgds (closed under fact tgd extension), it holds that:

PROPOSITION 6. $\mathsf{Eval}(\mathbb{O})$ *can be reduced in polynomial time into* $\mathsf{coCont}(\mathbb{O}, \mathbb{O}_\varnothing)$.

By definition, $\mathbb{O}_\varnothing$ is contained in every OMQ language $(\mathbb{C}, \mathbb{CQ})$, where $\mathbb{C}$ is a class of tgds. Therefore, as a corollary of Propositions 5 and 6, we obtain an initial boundary for the decidability of OMQ containment: we can obtain a positive result only if the evaluation problem for the involved OMQ languages is decidable. More formally:

COROLLARY 7. $\mathsf{Cont}(\mathbb{O}_1, \mathbb{O}_2)$ *is undecidable if* $\mathsf{Eval}(\mathbb{O}_1)$ *is undecidable or* $\mathsf{Eval}(\mathbb{O}_2)$ *is undecidable.*

Can we prove the converse of Corollary 7: $\mathsf{Cont}(\mathbb{O}_1, \mathbb{O}_2)$ is decidable if both $\mathsf{Eval}(\mathbb{O}_1)$ and $\mathsf{Eval}(\mathbb{O}_2)$ are decidable? The answer to this question is negative. This is due to the fact that containment of Datalog queries is undecidable [55]. Since Datalog queries can be directly encoded in the OMQ language based on the class $\mathbb{F}$ of *full tgds*, i.e., those without existentially quantified variables, we obtain the following:

PROPOSITION 8. [55] $\mathsf{Cont}((\mathbb{F}, \mathbb{CQ}))$ *is undecidable.*

This result, combined with the fact that $\mathsf{Eval}(\mathbb{F})$ is decidable (since the chase under full tgds always terminates), implies that the converse of Corollary 7 does not hold. Proposition 8 also rules out the OMQ languages that are based on classes of tgds that extend $\mathbb{F}$; e.g., the weak versions of the ones introduced in Section 2, called *weakly guarded* [24], *weakly acyclic* [35], and *weakly sticky* [27] that guarantee the decidability of OMQ evaluation.[2] The question that comes up concerns the decidability and complexity of containment for the OMQ languages that are based on the non-weak versions of the above classes, i.e., guarded, non-recursive, and sticky. This will be the subject of the next two sections.

3.2 From UCQs to CQs

Before we proceed with the complexity analysis of containment for the OMQ languages in question, let us state the following useful result:

PROPOSITION 9. *Given an OMQ* $Q \in (\mathbb{C}, \mathbb{UCQ})$*, where* $\mathbb{C} \in \{\mathbb{G}, \mathbb{L}, \mathbb{NR}, \mathbb{S}\}$*, we can construct in polynomial time an OMQ* $Q' \in (\mathbb{C}, \mathbb{CQ})$ *such that* $Q \equiv Q'$.

The proof of Proposition 9 relies on the idea of encoding boolean operations (in our case the 'or' operator) using a set of atoms; this idea has been used in several other works (see, e.g., [14, 21, 41]). Proposition 9 allows us to focus on OMQs that are based on CQs; in fact, $\mathsf{Cont}((\mathbb{C}_1, \mathbb{CQ}), (\mathbb{C}_2, \mathbb{CQ}))$ is $\mathcal{C}$-complete, where $\mathbb{C}_1, \mathbb{C}_2 \in \{\mathbb{G}, \mathbb{L}, \mathbb{NR}, \mathbb{S}\}$ and $\mathcal{C}$ is a complexity class that is closed under polynomial time reductions, iff $\mathsf{Cont}((\mathbb{C}_1, \mathbb{UCQ}), (\mathbb{C}_2, \mathbb{UCQ}))$ is $\mathcal{C}$-complete.

[2]The idea of those classes is the same: relax the conditions in the definition of the class, so that only those positions that receive null values during the chase are taken into account.

3.3 Plan of Attack

We are now ready to proceed with the complexity analysis of containment for the OMQ languages in question. Our plan of attack can be summarized as follows:

- We consider, in Section 4, $\mathsf{Cont}((\mathbb{C}, \mathbb{CQ}))$, for $\mathbb{C} \in \{\mathbb{L}, \mathbb{NR}, \mathbb{S}\}$. These languages enjoy a crucial property, called UCQ rewritability, which is very useful for our purposes. This property allows us to show the following result: if the containment does not hold, then this is witnessed via a "small" database, which in turn allows us to devise simple guess-and-check algorithms.
- We then proceed, in Section 5, with $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}))$. This OMQ language does not enjoy UCQ rewritability, and the task of establishing a small witness property as above turned out to be challenging. However, we show the following: if the containment does not hold, then this is witnessed via a "tree-shaped" database, which allows us to devise a decision procedure based on two-way alternating parity automata on finite trees.
- In Section 6, we study the case where the OMQ containment problem involves two different languages. If the left-hand side language is UCQ rewritable, then we can devise a guess-and-check algorithm by exploiting the above small witness property. The challenging case is when the left-hand side language is $(\mathbb{G}, \mathbb{CQ})$, where again we employ techniques based on tree automata.

4. UCQ REWRITABLE LANGUAGES

We now focus on OMQ languages that enjoy the crucial property of UCQ rewritability.

Definition 1. (**UCQ Rewritability**) An OMQ language $(\mathbb{C}, \mathbb{CQ})$, where $\mathbb{C} \subseteq \mathbb{TGD}$, is *UCQ rewritable* if, for each OMQ $Q = (\mathbf{S}, \Sigma, q(\bar{x})) \in (\mathbb{C}, \mathbb{CQ})$ we can construct a UCQ $q'(\bar{x})$ such that $Q(D) = q'(D)$ for every $\mathbf{S}$-database D. ■

We proceed to establish our desired small witness property, based on UCQ rewritability. By the definition of UCQ rewritability, for each language $\mathbb{O}$ that is UCQ rewritable, there exists a computable function $f_\mathbb{O}$ from $\mathbb{O}$ to the natural numbers such that the following holds: for every OMQ $Q = (\mathbf{S}, \Sigma, q(\bar{x})) \in \mathbb{O}$, and UCQ rewriting $q_1(\bar{x}) \vee \cdots \vee q_n(\bar{x})$ of Q, it is the case that $\max_{1 \leq i \leq n}\{|q_i|\} \leq f_\mathbb{O}(Q)$, where $|q_i|$ denotes the number of atoms occurring in q_i. Then:

PROPOSITION 10. *Consider a UCQ rewritable language* $\mathbb{O}$*, and two OMQs* $Q \in \mathbb{O}$ *and* $Q' \in (\mathbb{TGD}, \mathbb{CQ})$*, both with data schema* $\mathbf{S}$*. If* $Q \not\subseteq Q'$*, then there exists an* $\mathbf{S}$*-database* D*, where* $|D| \leq f_\mathbb{O}(Q)$*, such that* $Q(D) \not\subseteq Q'(D)$.

In Proposition 10 we assume that the left-hand side query falls in a UCQ rewritable language, be we do not pose any restriction on the language of the right-hand side query. Thus, we immediately get a decision procedure for $\mathsf{Cont}(\mathbb{O}_1, \mathbb{O}_2)$ if $\mathbb{O}_1$ is UCQ rewritable and $\mathsf{Eval}(\mathbb{O}_2)$ is decidable. Given $Q_1 = (\mathbf{S}, \Sigma_1, q_1(\bar{x})) \in \mathbb{O}_1$ and $Q_2 = (\mathbf{S}, \Sigma_2, q_2(\bar{x})) \in \mathbb{O}_2$:

1. Guess an $\mathbf{S}$-database D such that $|D| \leq f_{\mathbb{O}_1}(Q_1)$, and a tuple $\bar{c} \in dom(D)^{|\bar{x}|}$; and
2. Verify that $\bar{c} \in Q_1(D)$ and $\bar{c} \notin Q_2(D)$.

We immediately get that:

THEOREM 11. $\mathsf{Cont}(\mathbb{O}_1, \mathbb{O}_2)$ *is decidable if* $\mathbb{O}_1$ *is UCQ rewritable and* $\mathsf{Eval}(\mathbb{O}_2)$ *is decidable.*

This general result shows that $\mathsf{Cont}((\mathbb{C}, \mathbb{CQ}))$ is decidable for every $\mathbb{C} \in \{\mathbb{L}, \mathbb{NR}, \mathbb{S}\}$, but it says nothing about its complexity. This will be the subject of the rest of the section.

4.1 Linearity

The problem of computing UCQ rewritings for OMQs in $(\mathbb{L}, \mathbb{CQ})$ has been studied in [40], where a resolution-based procedure, called XRewrite, has been proposed. This rewriting algorithm accepts a query $Q = (\mathbf{S}, \Sigma, q(\bar{x})) \in (\mathbb{L}, \mathbb{CQ})$ and constructs a UCQ rewriting $q'(\bar{x})$ over $\mathbf{S}$ by starting from q and exhaustively applying rewriting steps based on resolution. Let us illustrate this via a simple example:

Example 1. Assume that $\mathbf{S} = \{P, T\}$. Consider the set Σ consisting of the linear tgds

$$P(x) \rightarrow \exists y\, R(x, y), \qquad R(x, y) \rightarrow P(y), \qquad T(x) \rightarrow P(x),$$

and the CQ $q(\bar{x}) \coloneqq \exists y(R(x, y) \wedge P(y))$. XRewrite will first resolve the atom $P(y)$ in q using the second tgd, and produce the CQ $\exists y(R(x, y) \wedge R(x, z))$, which is equivalent to the CQ $\exists y\, R(x, y)$. Then, $\exists y\, R(x, y)$ will be resolved using the first tgd, and the CQ $P(x)$ will be obtained, which in turn will be resolved using the third tgd in order to produce the CQ $T(x)$. The UCQ rewriting $q'(\bar{x})$ is $P(x) \vee T(x)$. ∎

It is easy to see that, whenever the input OMQ consists of linear tgds, during the execution of XRewrite it is not possible to obtain a CQ that has more atoms than the original one. This is an immediate consequence of the fact that linear tgds have only one atom in their body. Then:

PROPOSITION 12. $f_{(\mathbb{L}, \mathbb{CQ})}\big((\mathbf{S}, \Sigma, q)\big) \;\leq\; |q|$.

Having the above result in place, it can be shown that the algorithm underlying Theorem 11 guesses a polynomially sized witness to non-containment, and then calls a $\mathcal{C}$-oracle for solving query evaluation under linear OMQs, where $\mathcal{C}$ is PSPACE in general, and NP if the arity is fixed; these complexity classes are obtained from Proposition 2. Therefore, $\mathsf{coCont}((\mathbb{L}, \mathbb{CQ}))$ is in PSPACE in general, and in Σ_2^P in case of fixed arity. Regarding the lower bounds, Proposition 5 allows us to inherit the PSPACE-hardness of $\mathsf{Eval}(\mathbb{L}, \mathbb{CQ})$; this holds even for constant-free tgds. Unfortunately, in the case of fixed arity, we can only obtain NP-hardness, while Proposition 6 allows to obtain coNP-hardness. Nevertheless, it is implicit in [17] (see the proof of Theorem 9), where the containment problem for OMQ languages based on description logics is considered, that $\mathsf{Cont}((\mathbb{L}, \mathbb{CQ}))$ is Π_2^P-hard, even for tgds of the form $P(x) \rightarrow R(x)$. Then:

THEOREM 13. $\mathsf{Cont}((\mathbb{L}, \mathbb{CQ}))$ *is* PSPACE*-complete, and* Π_2^P*-complete if the arity of the schema is fixed. The lower bounds hold even for tgds without constants.*

4.2 Non-Recursiveness

Although the OMQ language $(\mathbb{NR}, \mathbb{CQ})$ is not explicitly considered in [40], where the algorithm XRewrite is defined, the same algorithm can deal with $(\mathbb{NR}, \mathbb{CQ})$. By analyzing the UCQ rewritings constructed by XRewrite, whenever the input query falls in $(\mathbb{NR}, \mathbb{CQ})$, we can establish the following result; here, $\mathit{body}(\tau)$ denotes the body of the tgd τ:

PROPOSITION 14. *It holds that*

$$f_{(\mathbb{NR}, \mathbb{CQ})}\big((\mathbf{S}, \Sigma, q)\big) \;\leq\; |q| \cdot \left(\max_{\tau \in \Sigma}\{|\mathit{body}(\tau)|\}\right)^{|sch(\Sigma)|}.$$

Proposition 14 implies that non-containment for queries that fall in $(\mathbb{NR}, \mathbb{CQ})$ is witnessed via a database of at most exponential size. We show next that this bound is optimal:

PROPOSITION 15. *There are sets of* $(\mathbb{NR}, \mathbb{CQ})$ *OMQs*

$$\{Q_1^n = (\mathbf{S}, \Sigma_1^n, q_1)\}_{n>0} \quad \textit{and} \quad \{Q_2^n = (\mathbf{S}, \Sigma_2^n, q_2)\}_{n>0},$$

where $|sch(\Sigma_1^n)| = |sch(\Sigma_2^n)| = n + 2$*, such that for every* $\mathbf{S}$*-database* D*, if* $Q_1^n(D) \not\subseteq Q_2^n(D)$ *then* $|D| \geq 2^{n-1}$.

Let us now focus on the complexity of $\mathsf{Cont}((\mathbb{NR}, \mathbb{CQ}))$. The algorithm underlying Theorem 11, together with the exponential bound provided by Proposition 14, implies that $\mathsf{coCont}((\mathbb{NR}, \mathbb{CQ}))$ is feasible in non-deterministic exponential time with access to a NEXPTIME oracle, which immediately implies that $\mathsf{Cont}((\mathbb{NR}, \mathbb{CQ}))$ is in EXPSPACE. Unfortunately, the exact complexity of $\mathsf{Cont}((\mathbb{NR}, \mathbb{CQ}))$ is still an open problem, and we conjecture that is $\mathrm{P}^{\mathrm{NEXP}}$-complete; recall that $\mathrm{NEXPTIME} \subseteq \mathrm{P}^{\mathrm{NEXP}} \subseteq \mathrm{EXPSPACE}$. In what follows, we briefly explain how the $\mathrm{P}^{\mathrm{NEXP}}$-hardness is obtained. To this end, we exploit a tiling problem that has been recently introduced in [34]. Roughly speaking, an instance of this tiling problem is a triple (m, T_1, T_2), where m is an integer in unary representation, and T_1, T_2 are standard tiling problems for the exponential grid $2^n \times 2^n$. The question is whether, for every initial condition w of length m, T_1 has no solution with w or T_2 has some solution with w. The initial condition w simply fixes the first m tiles of the first row of the grid. We construct in polynomial time two $(\mathbb{NR}, \mathbb{CQ})$ queries Q_1 and Q_2 such that (m, T_1, T_2) has a solution iff $Q_1 \subseteq Q_2$. The idea is to force every input database to store an initial condition w of length m, and then encode the problem whether T_i has a solution with w into Q_i, for each $i \in \{1, 2\}$. From the above discussion we get that:

THEOREM 16. $\mathsf{Cont}((\mathbb{NR}, \mathbb{CQ}))$ *is in* EXPSPACE*, and* $\mathrm{P}^{\mathrm{NEXP}}$*-hard. The lower bound holds even if the arity of the schema is fixed and the tgds are without constants.*

4.3 Stickiness

We now focus on $(\mathbb{S}, \mathbb{CQ})$. As shown in [40], given a query $(\mathbf{S}, \Sigma, q)$, there exists an execution of XRewrite that constructs a UCQ rewriting $q_1(\bar{x}) \vee \cdots \vee q_n(\bar{x})$ over $\mathbf{S}$ with the following property: for each $i \in \{1, \ldots, n\}$, if a variable v occurs in q_i in more than one atom, then v already occurs in q. This property has been used in [40] to bound the number of atoms that can appear in a single CQ q_i. We write $T(q)$ for the set of terms (constants and variables) occurring in q, $C(\Sigma)$ for the set of constants occurring in Σ, and $ar(\mathbf{S})$ for the maximum arity over all predicates of $\mathbf{S}$.

PROPOSITION 17. *It holds that*

$$f_{(\mathbb{S}, \mathbb{CQ})}((\mathbf{S}, \Sigma, q)) \;\leq\; |\mathbf{S}| \cdot (|T(q)| + |C(\Sigma)| + 1)^{|ar(\mathbf{S})|}.$$

Proposition 17 implies that non-containment for $(\mathbb{S}, \mathbb{CQ})$ queries is witnessed via a database of at most exponential size. As for $(\mathbb{NR}, \mathbb{CQ})$ queries, we can show that this bound is optimal; here, for a set Σ of tgds, we denote by $||\Sigma||$ the number of symbols occurring in Σ:

PROPOSITION 18. *There exists a set of* $(\mathbb{S}, \mathbb{CQ})$ *OMQs*

$$\{Q^n = (\{S/n\}, \Sigma^n, q(\bar{x}))\}_{n>0}, \text{ where } ||\Sigma^n|| \in O(n^2),$$

such that for every $Q = (\{S\}, \Sigma', q'(\bar{x})) \in (\mathbb{TGD}, \mathbb{CQ})$ *and* $\{S\}$*-database* D*, if* $Q^n(D) \not\subseteq Q(D)$ *then* $|D| \geq 2^{n-2}$.

We now study the complexity of $\mathsf{Cont}((\mathbb{S}, \mathbb{CQ}))$. We first focus on schemas of unbounded arity. Proposition 17 implies that the algorithm underlying Theorem 11 runs in exponential time assuming access to a $\mathcal{C}$-oracle, where $\mathcal{C}$ is a complexity class powerful enough for solving $\mathsf{Eval}(\mathbb{S}, \mathbb{CQ})$ and its complement. But, since $\mathsf{Eval}(\mathbb{S}, \mathbb{CQ})$ is in EXPTIME (see Proposition 4), both $\mathsf{Eval}(\mathbb{S}, \mathbb{CQ})$ and its complement are in NEXPTIME, and thus, the oracle call is not really needed. Consequently, $\mathsf{coCont}((\mathbb{C}, \mathbb{CQ}))$ is in NEXPTIME.

A matching lower bound is obtained by a reduction from the standard tiling problem for the exponential grid $2^n \times 2^n$. In fact, the same lower bound has been recently established in [15]; however, our result is stronger as it shows that the problem remains hard even if the right-hand side query is a linear OMQ of a simple form – this is also discussed in Section 6, where containment of queries that fall in different OMQ languages is studied. Regarding schemas of fixed arity, Proposition 17 provides a witness for non-containment of polynomial size, which implies that the algorithm underlying Theorem 11 runs in polynomial time with access to an NP-oracle. Therefore, $\mathsf{coEval}(\mathbb{S}, \mathbb{CQ})$ is in Σ_2^P, while a matching lower bound is implicit in [17]. Then:

THEOREM 19. $\mathsf{Cont}((\mathbb{S}, \mathbb{CQ}))$ *is* coNEXPTIME*-complete, even if the set of tgds uses only two constants. In the case of fixed arity, it is* Π_2^P*-complete, even for constant-free tgds.*

Clearly, there exists a double-exponential time algorithm for solving $\mathsf{Cont}((\mathbb{S}, \mathbb{CQ}))$, which might sound discouraging. However, Proposition 17 implies that the runtime is double-exponential only in the maximum arity of the data schema.

5. GUARDEDNESS

We proceed with the problem of containment for guarded OMQs, and we establish the following result:

THEOREM 20. $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}))$ *is* 2EXPTIME*-complete. The lower bound holds even if the arity of the schema is fixed, and the tgds are without constants.*

The lower bound is immediately inherited from [16], where it is shown that containment for OMQs based on the description logic $\mathcal{ELI}$ is 2EXPTIME-hard. Recall that a set of $\mathcal{ELI}$ axioms can be equivalently rewritten as a constant-free set of guarded tgds using only unary and binary predicates, which implies the lower bound stated in Theorem 20. However, we cannot immediately inherit the desired upper bound since the DL-based OMQ languages considered in [16] are either weaker than or incomparable to $(\mathbb{G}, \mathbb{CQ})$. Nevertheless, the technique developed in [16] was extremely useful for our analysis. Actually, our automata-based procedure exploits a combination of ideas from [16, 44]. The rest of this section is devoted to providing a high-level explanation of this procedure.

For the sake of technical clarity, we focus on constant-free tgds and CQs, but all the results can be extended to the general case at the price of more involved definitions and proofs. Moreover, for simplicity, we focus on Boolean CQs. In other words, we study the problem for $(\mathbb{G}, \mathbb{BCQ})$, where $\mathbb{BCQ}$ denotes the class of Boolean CQs. This does not affect the generality of our proof since it is known that $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}))$ can be reduced in polynomial time to $\mathsf{Cont}((\mathbb{G}, \mathbb{BCQ}))$ [16].

A first glimpse. As already said, $(\mathbb{G}, \mathbb{CQ})$ is not UCQ rewritable and, therefore, we cannot employ Proposition 10 in order to establish a small witness property as for the languages considered in Section 4. We have tried to establish a small witness property for $(\mathbb{G}, \mathbb{CQ})$ by following a different route, but it turned out to be a difficult task. Nevertheless, we can show a tree witness property, which states that non-containment for $(\mathbb{G}, \mathbb{CQ})$ is witnessed via a tree-like database. This allows us to devise a procedure based on alternating tree automata. Summing up, the proof for the 2EXPTIME membership of $(\mathbb{G}, \mathbb{CQ})$ proceeds in three steps:

1. Establish a tree witness property;
2. Encode the tree-like witnesses as trees that can be accepted by an alternating tree automaton; and
3. Construct an automaton that decides $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}))$; in fact, we reduce $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}))$ into emptiness for two-way alternating parity automata on finite trees.

Each one of the above three steps is discussed in more details in the following three sections.

5.1 Tree Witness Property

From the above informal discussion, it is clear that tree-like databases are crucial for our analysis. Let us make this notion more precise using guarded tree decompositions. A *tree decomposition* of a database D is a labeled rooted tree $T = (V, E, \lambda)$, where $\lambda : V \to 2^{dom(D)}$, such that: (i) for each atom $R(t_1, \ldots, t_n) \in D$, there exists $v \in V$ such that $\lambda(v) \supseteq \{t_1, \ldots, t_n\}$, and (ii) for every term $t \in dom(D)$, the set $\{v \in V \mid t \in \lambda(v)\}$ induces a connected subtree of T. The tree decomposition T is called $[U]$*-guarded*, where $U \subseteq V$, if, for every node $v \in V \setminus U$, there exists an atom $R(t_1, \ldots, t_n) \in D$ such that $\lambda(v) \subseteq \{t_1, \ldots, t_n\}$. We write $root(T)$ for the root node of T, and $D_T(v)$, where $v \in V$, for the subset of D induced by $\lambda(v)$. We are now ready to formalize the notion of the tree-like database:

Definition 2. An $\mathbf{S}$-database D is a C*-tree*, where $C \subseteq D$, if there is a tree decomposition T of D such that:

1. $D_T(root(T)) = C$ and
2. T is $[\{root(T)\}]$-guarded. ■

Roughly, whenever a database D is a C-tree, C is the cyclic part of D, while the rest of D is tree-like. Interestingly, for deciding $\mathsf{Cont}((\mathbb{G}, \mathbb{BCQ}))$ it suffices to focus on databases that are C-trees and $|dom(C)|$ depends only on the left-hand side OMQ. Recall that for a schema $\mathbf{S}$ we write $ar(\mathbf{S})$ for the maximum arity over all predicates of $\mathbf{S}$. Then:

PROPOSITION 21. *Let* $Q_i = (\mathbf{S}, \Sigma_i, q_i) \in (\mathbb{G}, \mathbb{BCQ})$*, for* $i \in \{1, 2\}$*. The following are equivalent:*

1. $Q_1 \subseteq Q_2$.
2. $Q_1(D) \subseteq Q_2(D)$*, for every* C*-tree* $\mathbf{S}$*-database* D *such that* $|dom(C)| \leq (ar(\mathbf{S} \cup sch(\Sigma_1)) \cdot |q_1|)$.

The fact that (1) ⇒ (2) holds trivially, while (2) ⇒ (1) is shown by using a variant of the notion of guarded unravelling and compactness. Let us clarify that the above result does not provide a decision procedure for $\mathsf{Cont}((\mathbb{G}, \mathbb{BCQ}))$, since we have to consider infinitely many databases that are C-trees with $|dom(C)| \leq (ar(\mathbf{S} \cup sch(\Sigma_1)) \cdot |q_1|)$.

5.2 Encoding Tree-like Databases

It is generally known that a database D whose treewidth[3] is bounded by an integer k can be encoded into a tree over a finite alphabet of double-exponential size in k that can be accepted by an alternating tree automaton; see, e.g., [13]. Consider an alphabet Γ, and let $\mathbb{N}^*$ be the set of finite sequences of natural numbers, including the empty sequence. A Γ-*labeled tree* is a pair $L = (T, \lambda)$, where $T \subseteq \mathbb{N}^*$ is closed under prefixes, and $\lambda \colon T \to \Gamma$ is the labeling function. The elements of T identify the nodes of L. It can be shown that D and a tree decomposition T of D with width k can be encoded as a Γ-labeled tree L, where Γ is an alphabet of double-exponential size in k, such that each node of T corresponds to exactly one node of L and vice versa.

Consider now a C-tree $\mathbf{S}$-database D, and let T be the tree decomposition that witnesses that D is a C-tree. The width of T is at most $k = (|dom(C)| + ar(\mathbf{S}) - 1)$, and thus, the treewidth of D is bounded by k. Hence, from the above discussion, D and T can be encoded as a Γ-labeled tree, where Γ is of double-exponential size in k. In general, given an $\mathbf{S}$-database D that is a C-tree due to the tree decomposition T, we show that D and T can be encoded as a $\Gamma_{\mathbf{S},l}$-labeled tree, with $|dom(C)| \leq l$ and $|\Gamma_{\mathbf{S},l}|$ being double-exponential in $ar(\mathbf{S})$ and exponential in $|\mathbf{S}|$ and l.

Although every C-tree $\mathbf{S}$-database D can be encoded as a $\Gamma_{\mathbf{S},l}$-labeled tree, the other direction does not hold. In other words, it is not true that every $\Gamma_{\mathbf{S},l}$-labeled tree encodes a C-tree $\mathbf{S}$-database D and its corresponding tree decomposition. In view of this fact, we need the additional notion of consistency. A $\Gamma_{\mathbf{S},l}$-labeled tree is called *consistent* if it satisfies certain syntactic properties – we do not give these properties here since they are not vital in order to understand the high-level idea of the proof. Now, given a consistent $\Gamma_{\mathbf{S},l}$-labeled tree L, we can show that L can be decoded into an $\mathbf{S}$-database $[\![L]\!]$ that is a C-tree with $|dom(C)| \leq l$. From the above discussion and Proposition 21, we obtain:

LEMMA 22. *Let $Q_i = (\mathbf{S}, \Sigma_i, q_i) \in (\mathbb{G}, \mathbb{BCQ})$, for $i \in \{1, 2\}$. The following are equivalent:*

1. *$Q_1 \subseteq Q_2$.*
2. *$Q_1([\![L]\!]) \subseteq Q_2([\![L]\!])$, for every consistent $\Gamma_{\mathbf{S},l}$-labeled tree L, where $l = (ar(\mathbf{S} \cup sch(\Sigma_1)) \cdot |q_1|)$.*

5.3 Constructing Tree Automata

Having the above result in place, we can now proceed with our automata-based procedure. We make use of two-way alternating parity automata (2WAPA) that run on finite labeled trees. Two-way alternating automata process the input tree while branching in an alternating fashion to successor states, and thereby moving either down or up the input tree; the detailed definition can be found in [11]. Our goal is to reduce $\mathsf{Cont}((\mathbb{G}, \mathbb{BCQ}))$ to the emptiness problem for 2WAPA. As usual, given a 2WAPA $\mathfrak{A}$, we denote by $\mathcal{L}(\mathfrak{A})$ the *language* of $\mathfrak{A}$, i.e., the set of labeled trees it accepts. The emptiness problem is defined as follows: given a 2WAPA $\mathfrak{A}$, does $\mathcal{L}(\mathfrak{A}) = \varnothing$? Thus, given $Q_1, Q_2 \in (\mathbb{G}, \mathbb{BCQ})$, we need to construct a 2WAPA $\mathfrak{A}$ such that $Q_1 \subseteq Q_2$ iff $\mathcal{L}(\mathfrak{A}) = \varnothing$. It is well-known that deciding whether $\mathcal{L}(\mathfrak{A}) = \varnothing$ is feasible in exponential time in the number of states, and in polynomial time in the size of the input alphabet [32]. Therefore, we should construct $\mathfrak{A}$ in double-exponential time, while the number of states must be at most exponential.

We first need a way to check consistency of labeled trees. It is not difficult to devise an automaton for this task.

LEMMA 23. *Consider a schema $\mathbf{S}$ and an integer $l > 0$. There is a 2WAPA $\mathfrak{C}_{\mathbf{S},l}$ that accepts a $\Gamma_{\mathbf{S},l}$-labeled tree L iff L is consistent. The number of states of $\mathfrak{C}_{\mathbf{S},l}$ is logarithmic in the size of $\Gamma_{\mathbf{S},l}$. Furthermore, $\mathfrak{C}_{\mathbf{S},l}$ can be constructed in polynomial time in the size of $\Gamma_{\mathbf{S},l}$.*

Now, the crucial task is, given an OMQ $Q \in (\mathbb{G}, \mathbb{BCQ})$, to devise an automaton that accepts labeled trees which correspond to databases that make Q true.

LEMMA 24. *Let $Q = (\mathbf{S}, \Sigma, q) \in (\mathbb{G}, \mathbb{BCQ})$. There is a 2WAPA $\mathfrak{A}_{Q,l}$, where $l > 0$, that accepts a consistent $\Gamma_{\mathbf{S},l}$-labeled tree L iff $Q([\![L]\!]) \neq \varnothing$. The number of states of $\mathfrak{A}_{Q,l}$ is exponential in $||Q||$ and l. Furthermore, $\mathfrak{A}_{Q,l}$ can be constructed in double-exponential time in $||Q||$ and l.*

The intuition underlying $\mathfrak{A}_{Q,l}$ can be described as follows. $\mathfrak{A}_{Q,l}$ tries to identify all the possible ways the CQ q can be mapped to $chase(D, \Sigma)$, for *any* C-tree $\mathbf{S}$-database D such that $|dom(C)| \leq l$. It then arrives at possible ways how the input tree can satisfy Q. These "possible ways" correspond to *squid decompositions*, a notion introduced in [24] that indicates which part of the query is mapped to the cyclic part C of D, and which to the tree-like part of D. The automaton exhaustively checks all squid decompositions by traversing the input tree and, at the same time, explores possible ways how to match the single parts of the squid decomposition at hand. The automaton finally accepts if it finds a squid decomposition that can be mapped to $chase(D, \Sigma)$.

Having the above automata in place, we can proceed with our main technical result, which shows that $\mathsf{Cont}(\mathbb{G}, \mathbb{BCQ})$ can be reduced to the emptiness problem for 2WAPA. But let us first recall some key results about 2WAPA, which are essential for our final construction. It is well-known that languages accepted by 2WAPAs are closed under intersection and complement. Given two 2WAPAs $\mathfrak{A}_1$ and $\mathfrak{A}_2$, we write $\mathfrak{A}_1 \cap \mathfrak{A}_2$ for a 2WAPA, which can be constructed in polynomial time, that accepts the language $\mathcal{L}(\mathfrak{A}_1) \cap \mathcal{L}(\mathfrak{A}_2)$. Moreover, for a 2WAPA $\mathfrak{A}$, we write $\overline{\mathfrak{A}}$ for the 2WAPA, which is also constructible in polynomial time, that accepts the complement of $\mathcal{L}(\mathfrak{A})$. We can now show the following:

PROPOSITION 25. *Consider $Q_1, Q_2 \in (\mathbb{G}, \mathbb{BCQ})$. We can construct in double-exponential time a 2WAPA $\mathfrak{A}$, which has exponentially many states, such that*

$$Q_1 \subseteq Q_2 \iff \mathcal{L}(\mathfrak{A}) = \varnothing.$$

PROOF (SKETCH). Let $Q_i = (\mathbf{S}, \Sigma_i, q_i)$, for $i \in \{1, 2\}$, and $l = (ar(\mathbf{S} \cup sch(\Sigma_1)) \cdot |q_1|)$. Then $\mathfrak{A}$ is defined as $(\mathfrak{C}_{\mathbf{S},l} \cap \mathfrak{A}_{Q_1,l}) \cap \overline{\mathfrak{A}_{Q_2,l}}$. Since $\Gamma_{\mathbf{S},l}$ has double-exponential size, Lemmas 23 and 24 imply that $\mathfrak{A}$ can be constructed in

[3]Recall that the treewidth of a database D is the minimum width among all possible tree decompositions $T = (V, E, \lambda)$ of D, while the width of T is defined as $\max_{v \in V}\{|\lambda(v)|\} - 1$.

double-exponential time, while it has exponentially many states. Lemma 22 implies that $Q_1 \subseteq Q_2$ iff $\mathcal{L}(\mathfrak{A}) = \varnothing$. □

Proposition 25 implies that $\mathsf{Cont}((\mathbb{G}, \mathbb{BCQ}))$ is in 2ExpTime, and Theorem 20 follows. Thus, there exists a double-exponential time algorithm for solving $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}))$. Interestingly, the runtime is double-exponential only in the size of the CQs and the maximum arity of the schema. This can be obtained by a providing a more refined complexity analysis of the construction of the 2WAPA $\mathfrak{A}$ in Proposition 25.

6. COMBINING LANGUAGES

In the previous two sections, we studied the containment problem relative to a language $\mathbb{O}$, i.e., both OMQs fall in $\mathbb{O}$. However, it is natural to consider the version of the problem where the involved OMQs fall in different languages. This is the goal of this section. Our analysis proceeds by considering the two cases where the left-hand side (LHS) query falls in a UCQ rewritable OMQ language, or it is guarded.

6.1 The LHS Query is UCQ Rewritable

As an immediate corollary of Theorem 11 we obtain the following result: $\mathsf{Cont}((\mathbb{C}_1, \mathbb{CQ}), (\mathbb{C}_2, \mathbb{CQ}))$, for $\mathbb{C}_1 \neq \mathbb{C}_2$, $\mathbb{C}_1 \in \{\mathbb{L}, \mathbb{NR}, \mathbb{S}\}$ and $\mathbb{C}_2 \in \{\mathbb{L}, \mathbb{NR}, \mathbb{S}, \mathbb{G}\}$, is decidable. By exploiting the algorithm underlying Theorem 11, we establish optimal upper bounds for all the problems at hand with the only exception of $\mathsf{Cont}((\mathbb{S}, \mathbb{CQ}), (\mathbb{NR}, \mathbb{CQ}))$. For the latter, we obtain an ExpSpace upper bound, by providing a similar analysis as for $\mathsf{Cont}((\mathbb{NR}, \mathbb{CQ}))$, while a NExpTime lower bound is inherited from query evaluation by exploiting Proposition 5. It is rather tedious, and not very interesting from a technical point of view, to go through all the containment problems in question[4] and explain in details how the exact upper bounds are obtained; we leave this as an exercise to the interested reader.

Regarding the matching lower bounds, in most of the cases they are inherited from query evaluation or its complement by exploiting Propositions 5 and 6, respectively. There are, however, some exceptions:

- $\mathsf{Cont}((\mathbb{S}, \mathbb{CQ}), (\mathbb{L}, \mathbb{CQ}))$ in the case of unbounded arity, where the problem is coNExpTime-hard, even for sets of tgds that use only two constants. This is shown by a reduction from the standard tiling problem for the exponential grid $2^n \times 2^n$.

- $\mathsf{Cont}((\mathbb{L}, \mathbb{CQ}), (\mathbb{S}, \mathbb{CQ}))$ and $\mathsf{Cont}((\mathbb{S}, \mathbb{CQ}), (\mathbb{L}, \mathbb{CQ}))$ in the case of bounded arity, where both problems are Π_2^P-hard even for constant-free tgds; implicit in [17].

6.2 The LHS Query is Guarded

We proceed with the case where the LHS query is guarded, and we show the following result:

Theorem 26. $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}), (\mathbb{C}, \mathbb{CQ}))$ *is* $\mathcal{C}$*-complete:*

$$\mathcal{C} = \begin{cases} \text{2ExpTime}, & \mathbb{C} \in \{\mathbb{L}, \mathbb{S}\}, \\ \text{3ExpTime}, & \mathbb{C} = \mathbb{NR}. \end{cases}$$

[4]There are eighteen different cases obtained by considering all the possible pairs $(\mathbb{O}_1, \mathbb{O}_2)$ of OMQ languages, where $\mathbb{O}_1 \neq \mathbb{O}_2$ and $\mathbb{O}_1$ is UCQ rewritable, and the two cases whether the arity of the schema is fixed or not.

The lower bounds hold even if the arity of the schema is fixed. Moreover, for $\mathbb{C} = \mathbb{L}$ (resp., $\mathbb{C} \in \{\mathbb{NR}, \mathbb{S}\}$) it holds even for tgds with one constant (resp., without constants).

Upper bounds. The 2ExpTime membership when $\mathbb{C} = \mathbb{L}$ is an immediate corollary of Theorem 20. This is not true when $\mathbb{C} \in \{\mathbb{NR}, \mathbb{S}\}$ since the right-hand side query is not guarded. But in this case, since $(\mathbb{NR}, \mathbb{CQ})$ and $(\mathbb{S}, \mathbb{CQ})$ are UCQ rewritable, one can rewrite the right-hand side query as a UCQ, and then apply the machinery developed in Section 5 for solving $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}))$. More precisely, given OMQs $Q_1 \in (\mathbb{G}, \mathbb{CQ})$ and $Q_2 \in (\mathbb{C}, \mathbb{CQ})$, where $\mathbb{C} \in \{\mathbb{NR}, \mathbb{S}\}$, $Q_1 \subseteq Q_2$ iff $Q_1 \subseteq q$, where q is a UCQ rewriting of Q_2. Thus, an immediate decision procedure, which exploits the algorithm $\mathsf{XRewrite}$, is the following:

1. Let $q = \mathsf{XRewrite}(Q_2)$;
2. For each $q' \in q$: if $Q_1 \subseteq q'$, then proceed; otherwise, reject; and
3. Accept.

The above procedure runs in triple-exponential time. The first step is feasible in double-exponential time [40]. Now, for a single CQ $q' \in q$ (which is a guarded OMQ with an empty set of tgds) the check whether $Q_1 \subseteq q'$ can be done by using the machinery developed in Section 5, which reduces our problem to checking whether the language of a 2WAPA $\mathfrak{A}$ is empty. However, it should not be forgotten that q' is of exponential size, and thus, the automaton $\mathfrak{A}$ has double-exponentially many states. This in turn implies that checking whether $\mathcal{L}(\mathfrak{A}) = \varnothing$ is in 3ExpTime, as claimed.

Although the above algorithm establishes an optimal upper bound for non-recursive OMQs, a more refined analysis is needed for sticky OMQs. In fact, we need a more refined complexity analysis for the problem $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}), \mathbb{UCQ})$, that is, to decide whether a guarded OMQ is contained in a UCQ. To this end, we provide an automata construction different from the one employed in Section 5, which allows us to establish a refined complexity upper bound for the problem in question. Consider a $(\mathbb{G}, \mathbb{CQ})$ query Q, and a UCQ $q = q_1 \vee \cdots \vee q_n$. As usual, we write $||Q||$ and $||q_i||$ for the number of symbols that occur in Q and q_i, respectively, and we write $var_{\geq 2}(q_i)$ for the set of variables that appear in more than one atom of q_i. By exploiting our new automata-based procedure, we show that the problem of checking if $Q \subseteq q$ is feasible in double-exponential time in $(||Q|| + \max_{1 \leq i \leq n}\{|var_{\geq 2}(q_i)|\})$, exponential time in $\max_{1 \leq i \leq n}\{||q_i||\}$, and polynomial time in n.

This result allows us to show that the above procedure establishes 2ExpTime-membership when the right-hand side OMQ is sticky. But first we need to recall the following key properties of the UCQ rewriting $q = \mathsf{XRewrite}(Q_2)$, constructed during the first step of the algorithm:

1. q consists of double-exponentially many CQs,
2. each CQ of q is of exponential size, and
3. for each $q' \in q$, $var_{\geq 2}(q')$ is a subset of the variables of the original CQ that appears in Q_2.

By combining these key properties with the complexity analysis performed above, it is now straightforward to show that $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}), (\mathbb{S}, \mathbb{CQ}))$ is in 2ExpTime.

Lower Bounds. We establish matching lower bounds by refining techniques from [31], where it is shown that containment of Datalog in UCQ is 2ExpTime-complete, while containment of Datalog in non-recursive Datalog is 3ExpTime-complete; the lower bounds hold for fixed-arity predicates, and constant-free rules. Interestingly, the LHS query can be transformed into a Datalog query such that each rule has a body-atom that contains all the variables, i.e., is guarded. This is achieved by increasing the arity of some predicates in order to have enough positions for all the body-variables. However, for each rule, the number of unguarded variables that we need to guard is constant, and thus, the arity of the schema remains constant. We conclude that $\mathsf{Cont}((\mathbb{G},\mathbb{CQ}),(\mathbb{NR},\mathbb{CQ}))$ is 3ExpTime-hard. Moreover, containment of guarded OMQs in UCQs is 2ExpTime-hard, which in turn allows us to show, by exploiting the construction underlying Proposition 9, that $\mathsf{Cont}((\mathbb{G},\mathbb{CQ}),(\mathbb{L},\mathbb{CQ}))$ is 2ExpTime-hard, even if the set of linear tgds uses only one constant, while $\mathsf{Cont}((\mathbb{G},\mathbb{CQ}),(\mathbb{S},\mathbb{CQ}))$ is 2ExpTime-hard, even for tgds without constants.

7. APPLICATIONS

Interestingly, our results on $\mathsf{Cont}((\mathbb{G},\mathbb{CQ}))$ can be applied to other important static analysis tasks, in particular, distribution over components and UCQ rewritability. Each one of those tasks is considered in the following two sections.

7.1 Distribution Over Components

The notion of distribution over components has been introduced in [3], and it states that the answer to a query can be computed by parallelizing it over the (maximally connected) components of the input database. But let us first make precise what a component is. A set of atoms A is *connected* if for all $c, d \in dom(A)$, there exists a sequence $\alpha_1, \ldots, \alpha_n$ of atoms in A such that $c \in dom(\alpha_1)$, $d \in dom(\alpha_n)$, and $dom(\alpha_i) \cap dom(\alpha_{i+1}) \neq \varnothing$, for each $i \in \{1, \ldots, n-1\}$. We call $B \subseteq A$ a *component* of A if (i) B is connected, and (ii) for every $\alpha \in A \setminus B$, $B \cup \{\alpha\}$ is not connected.[5] Let $co(A)$ be the set of components of A. We are now ready to introduce the notion of distribution over components. Consider an OMQ $Q = (\mathbf{S}, \Sigma, q) \in (\mathbb{TGD},\mathbb{CQ})$. We say that Q *distributes over components* if $Q(D) = Q(D_1) \cup \cdots \cup Q(D_n)$, where $co(D) = \{D_1, \ldots, D_n\}$, for every $\mathbf{S}$-database D. In this case, $Q(D)$ can be computed without any communication over a network using a distribution where every computing node is assigned some of the components of the database, and every component is assigned to at least one computing node. In other words, Q can be evaluated in a distributed and coordination-free manner; for more details on coordination-free evaluation see [3, 4, 5]. Therefore, it would be quite beneficial if we can decide whether an OMQ distributes over components, and thus, we obtain the following interesting static analysis task:

PROBLEM :	$\mathsf{Dist}(\mathbb{C},\mathbb{CQ})$
INPUT :	An OMQ $Q \in (\mathbb{C},\mathbb{CQ})$.
QUESTION :	Does Q distributes over components?

The above problem has been studied in [15], where tight complexity bounds for $(\mathbb{L},\mathbb{CQ})$ and $(\mathbb{S},\mathbb{CQ})$ have been established. However, its exact complexity for guarded OMQs has been left open. Our results on containment for guarded OMQs allow us to close this problem. But first we need to recall a key result that semantically characterizes distribution over components. An OMQ Q with data schema $\mathbf{S}$ is *unsatisfiable* if there is no $\mathbf{S}$-database D such that $Q(D) \neq \varnothing$. Moreover, for a CQ q, we write $co(q)$ for its components. The next result has been shown in [15]:

Proposition 27. *Let $Q = (\mathbf{S}, \Sigma, q(\bar{x})) \in (\mathbb{G},\mathbb{CQ})$. The following are equivalent:*

1. *Q distributes over components.*
2. *Q is unsatisfiable or there exists $\hat{q}(\bar{x}) \in co(q)$ such that $(\mathbf{S}, \Sigma, \hat{q}(\bar{x})) \subseteq Q$.*

Checking unsatisfiability can be easily reduced to containment. Thus, the above result, together with Theorem 20, implies that $\mathsf{Dist}(\mathbb{G},\mathbb{CQ})$ is in 2ExpTime, while a matching lower bound is implicit in [15]. Then:

Theorem 28. $\mathsf{Dist}(\mathbb{G},\mathbb{CQ})$ *is* 2ExpTime*-complete.*

7.2 Deciding UCQ Rewritability

Query rewriting is a well-studied method for evaluating OMQs using standard database technology. The key idea is the following: given an OMQ $Q = (\mathbf{S}, \Sigma, q(\bar{x}))$, combine Σ and q into a new query $q_\Sigma(\bar{x})$, the so-called rewriting, which can then be evaluated over D yielding the same answer as Q over D, for *every* $\mathbf{S}$-database D. For this approach to be realistic, though, it is essential that the rewriting is expressed in a language that can be handled by standard database systems. The typical language that is considered in this setting is first-order (FO) queries [28]. Notice, however, that due to Rossman's Theorem [53], and the fact that OMQs are closed under homomorphisms, FO and UCQ rewritability coincide. Recall that some OMQ languages are UCQ rewritable, such as the ones based on linear, non-recursive and sticky sets of tgds, while others are not, e.g., guarded OMQs. For those languages $\mathbb{O}$ that are not UCQ rewritable, it is important to be able to check whether a query $Q \in \mathbb{O}$ can be rewritten as a UCQ, in which case we say that it is UCQ rewritable. This gives rise to the following fundamental static analysis task for an OMQ language $(\mathbb{C},\mathbb{CQ})$, where $\mathbb{C} \subseteq \mathbb{TGD}$:

PROBLEM :	$\mathsf{UCQRew}(\mathbb{C},\mathbb{CQ})$
INPUT :	An OMQ $Q \in (\mathbb{C},\mathbb{CQ})$.
QUESTION :	Is it the case that Q is UCQ rewritable?

Bienvenu et al. have recently carried out an in-depth study of the above problem for OMQ languages based on central Horn-DLs [16]. One of their main results is that the above problem for the OMQ language based on $\mathcal{ELHI}$, one of the most expressive members of the $\mathcal{EL}$-family of DLs, is 2ExpTime-complete. Interestingly, by adapting the tree automata techniques developed in Section 5, we can generalize the above result: deciding UCQ rewritability for the OMQ language based on guarded tgds over unary and binary relations is in 2ExpTime. Let $\mathbb{G}_2$ be the class of (finite) sets of guarded tgds over unary and binary relations. Then:

Theorem 29. $\mathsf{UCQRew}(\mathbb{G}_2,\mathbb{CQ})$ *is* 2ExpTime*-complete.*

[5] For technical clarity, the notion of component is defined only for sets of atoms that do not contain 0-ary atoms.

Since the lower bound is inherited from [16], we concentrate on the upper bound. As in Section 5, we can focus on BCQs, i.e., it suffices to show that $\mathsf{UCQRew}(\mathbb{G}_2, \mathbb{BCQ})$ is in 2ExpTime. Our proof proceeds in two steps:

1. We semantically characterize UCQ rewritability for queries in $(\mathbb{G}_2, \mathbb{CQ})$ in terms of a certain *boundedness* property for the set of C-trees defined in Section 5.
2. We extend the techniques developed in Section 5 and construct in double-exponential time a 2WAPA $\mathfrak{A}$ that has exponentially many states, such that the aforementioned boundedness property does not hold iff $\mathcal{L}(\mathfrak{A})$ is infinite. (Such an *infinity problem* for tree automata has been used to obtain the decidability of the *boundedness* problem for monadic Datalog [32, 56]).

Our 2ExpTime upper bound then follows since the infinity problem for a 2WAPA $\mathfrak{A}$, i.e., checking if $\mathcal{L}(\mathfrak{A})$ is infinite, is feasible in exponential time in the number of states, and in polynomial time in the size of the alphabet. This follows from two known results: (a) The 2WAPA $\mathfrak{A}$ can be converted into an equivalent non-deterministic tree automata $\mathfrak{B}$ with a single-exponential blow up in the number of states [57], and (b) solving the infinity problem for non-deterministic tree automata is feasible in polynomial time; cf. [56].

It is worth contrasting our proof with the one in [16] for $\mathcal{ELHI}$, which does not make use of the infinity problem for 2WAPA, but applies a different argument based on pumping. This leads to a finer complexity analysis in terms of the size of the different components of the OMQ, but, in our opinion, makes the proof conceptually harder.

The semantic characterization. To establish the semantic characterization from step 1, we need to define the notion of *distance from the root* for an element u in a C-tree database D. Intuitively, this corresponds to the minimal distance between a node that contains u and the root of a tree decomposition T of D that witnesses the fact that D is a C-tree. We do not consider all such tree decompositions, however, but concentrate on a well-behaved subclass, which we call the *lean* tree decompositions of the C-tree D; the formal definition can be found in [11], as it does not add much to the explanation we provide here. Due to the fact that we focus on unary and binary relations, such lean tree decompositions ensure the invariance of the notion of distance from the root, by severely limiting the level of redundancy allowed in a tree representation of D. Therefore, it does not matter which lean tree decomposition we choose, since in all of them the distance of an element u from the root will be the same. Let $D_{\leq k}$ be the subinstance of D induced by the set of elements whose distance from the root is at most k, and let $D_{>k}$ be the subinstance of D induced by the set of elements whose distance from the root is at least $k+1$.

Another useful notion is the *branching degree* of a tree decomposition T, that is, the maximum number of child nodes over all nodes of T. Again, lean tree decompositions ensure the invariance of the branching degree. This allows us to define the branching degree of a C-tree database D as the branching degree of a lean tree decomposition that witnesses the fact that D is a C-tree.

It follows from [16] that being able to decide containment for the OMQ language $(\mathbb{G}_2, \mathbb{BCQ})$ (as we have done in Section 5) allows us to concentrate on *connected* CQs when deciding UCQ rewritability. This simplifies technicalities considerably and, in turn, allows us to obtain our desired semantic characterization of UCQ rewritability:

Proposition 30. *Let $Q = (\mathbf{S}, \Sigma, q) \in (\mathbb{G}_2, \mathbb{BCQ})$, where q is connected. The following are equivalent:*

1. *Q is UCQ rewritable.*
2. *There exist $k, m \geq 0$ (which depend only on Q) s.t.*

$$Q(D) \neq \varnothing \implies \big(Q(D_{\leq k}) \neq \varnothing \text{ or } Q(D_{>0}) \neq \varnothing\big),$$

for each C-tree $\mathbf{S}$-database D with $|dom(C)| \leq 2 \cdot |q|$ and branching degree at most m.

The reduction to the infinity problem. We now proceed with step 2, and we explain how the boundedness property established in item (2) of Proposition 30 can be reduced to the infinity problem for 2WAPAs. As in Section 5, we do not reason with C-tree databases directly, but we deal with their encodings as consistent $\Gamma_{\mathbf{S},l}$-labeled trees. In fact, using the same ideas as in Lemma 22, we can show by exploiting Proposition 30 that the following are equivalent:

(i) Q is UCQ rewritable.

(ii) There are $k, m \geq 0$ such that

$$Q([\![L]\!]) \neq \varnothing \implies \big(Q([\![L]\!]_{\leq k}) \neq \varnothing \text{ or } Q([\![L]\!]_{>0}) \neq \varnothing\big),$$

for every consistent $\Gamma_{\mathbf{S},l}$-labeled tree L with $l = 2 \cdot |q|$ and whose branching degree is bounded by m.

Let us write Boundedness for the property expressed in item (ii) above, which can be reduced to the problem of checking whether some tree language is finite. Let $\mathcal{L}_Q$ be the set of all $\Gamma_{\mathbf{S},l}$-labeled trees L of branching degree at most m such that: (1) $Q([\![L]\!]) = \varnothing$ and (2) there is some "extension" L' of L, with branching degree m, such that $Q([\![L']\!]) \neq \varnothing$ and $Q([\![L']\!]_{>0}) = \varnothing$. Notice that L' can increase the depth but not the branching degree of L. It is not difficult to show that Boundedness holds iff $\mathcal{L}_Q$ is finite. We then devise in double-exponential time a 2WAPA $\mathfrak{C}_{Q,l}$, which has exponentially many states, such that $\mathcal{L}_Q = \mathcal{L}(\mathfrak{C}_{Q,l})$. Therefore, the following holds:

Proposition 31. *Consider $Q \in (\mathbb{G}_2, \mathbb{BCQ})$. We can construct in double-exponential time a 2WAPA $\mathfrak{A}$, which has exponentially many states, such that*

$$Q \text{ is UCQ rewritable} \iff \mathcal{L}(\mathfrak{A}) \text{ is finite.}$$

Since checking whether $\mathcal{L}(\mathfrak{A})$ is infinite is feasible in exponential time in the number of states and in polynomial time in the size of the alphabet, Proposition 31 implies that $\mathsf{UCQRew}(\mathbb{G}_2, \mathbb{CQ})$ is in 2ExpTime, as needed.

8. CONCLUSIONS

We have concentrated on the fundamental problem of containment for OMQ languages based on the main decidable classes of tgds. We have also used our techniques to close problems related to distribution over components and UCQ rewritability. We believe that our techniques for solving containment under guarded OMQs can be extended to frontier-guarded OMQs, an interesting extension of guardedness [8]. We are also convinced that our solution to the problem of deciding UCQ rewritability of guarded OMQs over unary and binary relations can be extended to guarded (or even frontier-guarded) OMQs over arbitrary schemas. We are currently investigating these challenging problems.

APPENDIX

PRELIMINARIES

Definition of Non-recursiveness

In the main body of the paper, we define non-recursive sets of tgds via the notion of predicate graph. Here, we give an alternative definition, based on the well-known notion of stratification, which is more convenient for the combinatorial analysis that we are going to perform in the proof of Proposition 14.

Definition 3. Consider a set Σ of tgds. A *stratification* of Σ is a partition $\{\Sigma_1, \ldots, \Sigma_n\}$, where $n > 0$, of Σ such that, for some function $\mu : sch(\Sigma) \to \{0, \ldots, n\}$, the following hold:

1. For each predicate $R \in sch(\Sigma)$, all the tgds with R in their head belong to $\Sigma_{\mu(R)}$, i.e., they belong to the same set of the partition.

2. If there exists a tgd in Σ such that the predicate R appears in its body, while the predicate P appears in its head, then $\mu(R) < \mu(P)$.

We say that Σ is *stratifiable* if it admits a stratification. ∎

It is an easy exercise to show that the predicate graph of a set Σ of tgds is acyclic iff Σ is stratifiable. Then:

LEMMA 32. *Σ is non-recursive iff Σ is stratifiable.*

Definition of Stickiness

In the main body of the paper, we provide an intuitive explanation of stickiness. Here, we recall the formal definition of sticky sets of tgds, introduced in [27]. Fix a set Σ of tgds; w.l.o.g., we assume that, for every pair $(\sigma, \sigma') \in \Sigma \times \Sigma$, σ and σ' do not share variables. For notational convenience, given an atom α and a variable x occurring in α, $pos(\alpha, x)$ is the set of positions in α at which x occurs; a position $P[i]$ identifies the i-th attribute of the predicate P. The definition of stickiness hinges on the notion of marked variables in a set of tgds.

Definition 4. Consider a tgd $\sigma \in \Sigma$, and a variable x occurring in the body of σ. We inductively define when x is *marked in* Σ as follows:

1. If there exists an atom α in the head of σ such that x does not occur in α, then x is marked in Σ; and

2. Assuming that there exists an atom α in the head of σ such that x occurs in α, if there exists $\sigma' \in \Sigma$ (not necessarily different than σ) and an atom β in the body of σ' such that (i) α and β have the same predicate and, (ii) each variable in β that occurs at a position of $pos(\alpha, x)$ is marked in Σ, then x is marked in Σ. ∎

We are now ready to recall when a set of tgds is sticky:

Definition 5. A set Σ of tgds is *sticky* if, for each $\sigma \in \Sigma$, and for each variable x occurring in the body of σ, the following holds: if x is marked in Σ, then x occurs only once in the body of σ. ∎

PROOFS OF SECTION 3

Proof of Proposition 5

Consider an OMQ $Q = (\mathbf{S}, \Sigma, q(\bar{x})) \in (\mathbb{C}, \mathbb{CQ})$, where $\mathbb{C}$ is a class of tgds, an $\mathbf{S}$-database D, and a tuple $\bar{c} \in dom(D)^{|\bar{x}|}$. We show that:

$$\bar{c} \in Q(D) \iff \underbrace{(sch(\Sigma), \varnothing, q_{D,\bar{c}})}_{Q_1} \subseteq \underbrace{(sch(\Sigma), \Sigma, q)}_{Q_2}.$$

($\Rightarrow$) Assume that $Q_1 \not\subseteq Q_2$. This implies that there exists a $sch(\Sigma)$-database D', and a tuple $\bar{t}$ of constants such that $\bar{t} \in q_{D,\bar{c}}(D')$ and $\bar{t} \notin q(chase(D', \Sigma))$. Due to the monotonicity of CQs, $\bar{t} \in q_{D,\bar{c}}(chase(D', \Sigma))$. Since, by construction, the instance $chase(D', \Sigma)$ satisfies Σ, we conclude that $q_{D,\bar{c}} \not\subseteq_\Sigma q$.[6] By exploiting the well-known characterization of CQ containment in terms of the chase, we get that $\bar{c} \notin q(chase(D, \Sigma))$, which is equivalent to $\bar{c} \notin Q(D)$, as needed.

($\Leftarrow$) Conversely, assume that $\bar{c} \notin Q(D)$, or, equivalently, $\bar{c} \notin q(chase(D, \Sigma))$. This implies that $\bar{c} \notin Q_2(D)$. Observe that $\bar{c} \in q_{D,\bar{c}}(D)$ holds trivially, which in turn implies that $\bar{c} \in Q_1(D)$. Therefore, $Q_1 \not\subseteq Q_2$, and the claim follows.

[6] This is the standard notation for the fact that $q_{D,\bar{c}}(I) \not\subseteq q(I)$, for every (possibly infinite) instance I that satisfies Σ.

Proof of Proposition 9

The construction underlying Proposition 9 relies on the idea of encoding boolean operations (in our case the 'or' operator) using a set of atoms; this idea has been exploited in several other works; see, e.g., [14, 21, 41]. Let $Q = (\mathbf{S}, \Sigma, q) \in (\mathbb{C}, \mathbb{UCQ})$. Our goal is to construct in polynomial time $Q' = (\mathbf{S}, \Sigma', q') \in (\mathbb{C}, \mathbb{CQ})$ such that $Q \equiv Q'$. We assume, w.l.o.g., that the predicates of $\mathbf{S}$ do not appear in the head of a tgd of Σ; we can copy the content of a relation $R/k \in \mathbf{S}$ into an auxiliary predicate $R^\star/k$, using the tgd $R(x_1, \ldots, x_k) \rightarrow R^\star(x_1, \ldots, x_k)$, while staying inside $\mathbb{C}$, and then rename each predicate P in Σ and q with $P^\star$. The set Σ' consists of the following tgds:

1. For every $R/k \in \mathbf{S}$:

$$R(x_1, \ldots, x_k) \ \rightarrow \ R'(x_1, \ldots, x_k, 1), \mathrm{True}(1).$$

   These tgds are annotating the database atoms with the truth constant true, indicating that these are true atoms.

2. Assuming that $q = \exists \bar{y}\, \phi(\bar{x}, \bar{y})$, a tgd:

$$\mathrm{True}(t) \ \rightarrow \exists \bar{x} \exists \bar{y} \exists f\, \phi'_\wedge(\bar{x}, \bar{y}, f), \psi(t, f),$$

   where $\phi'_\wedge$ is the conjunction of atoms in ϕ, after replacing each atom $R(v_1, \ldots, v_k)$ with $R'(v_1, \ldots, v_k, f)$, and ψ is the conjunction of atoms

$$\mathrm{Or}(t,t,t), \mathrm{Or}(t,f,t), \mathrm{Or}(f,t,t), \mathrm{Or}(f,f,f).$$

   This tgd generates a "copy" of the atoms in q, while annotating them with a null value that represents the truth constant false, indicating that are not necessarily true atoms. Moreover, the truth table of 'or' is generated.

3. Finally, for each tgd $\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z}\, \psi(\bar{x}, \bar{x})$ in Σ, a tgd

$$\phi'(\bar{x}, \bar{y}, w) \ \rightarrow \ \exists \bar{z}\, \psi'(\bar{x}, \bar{z}, w),$$

   where ϕ' and ψ' are obtained from ϕ and ψ, respectively, by replacing each atom $R(v_1, \ldots, v_k)$ with $R'(v_1, \ldots, v_k, w)$. In fact, this is the actual set of tgds Σ, with the difference that the value at the last position of each atom (which indicates whether it is true or false) is propagated to the inferred atoms.

Now, assuming that $q = q_1 \vee \cdots \vee q_n$, the CQ q' is defined as follows; let $\bar{x} = x_1 \ldots x_n$ and $\bar{y} = y_1 \ldots y_{n+1}$:

$$\exists \bar{x} \exists \bar{y}\, (\mathrm{False}(y_1) \ \wedge \bigwedge_{1 \leq i \leq n} (q'_i[x_i] \wedge \mathrm{Or}(y_i, x_i, y_{i+1})) \ \wedge \mathrm{True}(y_{n+1})),$$

where $\bar{x}$ and $\bar{y}$ are fresh variables not in q, and $q'_i[x_i]$ is obtained from q_i by replacing each atom $R(v_1, \ldots, v_k)$ with $R'(v_1, \ldots, v_k, x_i)$. This completes our construction.

It is not difficult to show that $Q \equiv Q'$, or, equivalently, for every $\mathbf{S}$-database D, $q(\mathit{chase}(D, \Sigma)) = q'(\mathit{chase}(D, \Sigma'))$. The key observation is that in order to satisfy $\mathrm{True}(y_{n+1})$ in the CQ q', at least one of the $\bar{x}_i$'s must be mapped to 1, which means that at least one q_i is satisfied by $\mathit{chase}(D, \Sigma)$. Finally, it is easy to verify that, for each $\mathbb{C} \in \{\mathbb{G}, \mathbb{L}, \mathbb{NR}, \mathbb{S}\}$, $\Sigma \in \mathbb{C}$ implies $\Sigma' \in \mathbb{C}$, and Proposition 9 follows.

PROOFS OF SECTION 4

Proof of Proposition 10

We assume that $q(\bar{x}) = \bigvee_{i=1}^n q_i(\bar{x})$ is a UCQ rewriting of Q. Since, by hypothesis, $Q \not\subseteq Q'$, we conclude that $q \not\subseteq Q'$, which in turn implies that there exists $i \in \{1, \ldots, n\}$ such that $q_i \not\subseteq Q'$. Let $c(\bar{x})$ be a tuple of constants obtained by replacing each variable x in $\bar{x}$ with the constant $c(x)$, and D_{q_i} the $\mathbf{S}$-database obtained from q_i after replacing each variable x in q_i with the constant $c(x)$. We show that:

LEMMA 33. $c(\bar{x}) \notin Q'(D_{q_i})$.

PROOF. Since $q_i \not\subseteq Q'$, there exists an $\mathbf{S}$-database D, and a tuple of constants $\bar{t}$ such that $\bar{t} \in q_i(D)$ and $\bar{t} \notin Q'(D)$. Clearly, there exists a homomorphism h such that $h(q_i) \subseteq D$ and $h(\bar{x}) = \bar{t}$. Observe also that $\rho(D_{q_i}) \subseteq D$, where $\rho = h \circ c^{-1}$. Towards a contradiction, assume that $c(\bar{x}) \in Q'(D_{q_i})$. This implies that there exists a homomorphism γ such that $\gamma(q') \subseteq \mathit{chase}(D_{q_i}, \Sigma)$ and $\gamma(\bar{y}) = c(\bar{x})$, where $Q' = (\mathbf{S}, \Sigma, q'(\bar{y}))$. It is not difficult to see that there exists an extension ρ' of ρ such that $\rho'(\mathit{chase}(D_{q_i}, \Sigma)) \subseteq \mathit{chase}(D, \Sigma)$ and $\rho'(\bar{x}) = \bar{t}$. Hence, $\rho'(\gamma(q')) \subseteq \mathit{chase}(D, \Sigma)$, which implies that $\bar{t} \in q'(\mathit{chase}(D, \Sigma))$; thus, $\bar{t} \in Q'(D)$. But this contradicts the fact that $\bar{t} \notin Q'(D)$, and the claim follows. □

Observe that $c(\bar{x}) \in q(D_{q_i})$, which immediately implies that $c(\bar{x}) \in Q(D_{q_i})$. Consequently, by Lemma 33, $Q(D_{q_i}) \not\subseteq Q'(D_{q_i})$. The claim follows since, by construction, D_{q_i} is an $\mathbf{S}$-database such that $|D_{q_i}| \leq f_{\mathbb{O}}(Q)$.

Algorithm 1: The algorithm XRewrite

Input: An OMQ $Q = (\mathbf{S}, \Sigma, q(\bar{x})) \in (\mathbb{TGD}, \mathbb{CQ})$
Output: A UCQ $q'(\bar{x})$ such that $Q(D) = q'(D)$, for every $\mathbf{S}$-database D

$i := 0$;
$Q_{\text{REW}} := \{\langle q, \mathsf{r}, \mathsf{u}\rangle\}$;
repeat
  $Q_{\text{TEMP}} := Q_{\text{REW}}$;
  foreach $\langle q, x, \mathsf{u}\rangle \in Q_{\text{TEMP}}$, where $x \in \{\mathsf{r}, \mathsf{f}\}$ **do**
    foreach $\sigma \in \Sigma$ **do**
      /* rewriting step */
      foreach $S \subseteq \mathit{body}(q)$ such that σ is applicable to S **do**
        $i := i + 1$;
        $q' := \gamma_{S,\sigma^i}(q[S/\mathit{body}(\sigma^i)])$;
        if there is no $(q'', \mathsf{r}, \star) \in Q_{\text{REW}}$ such that $q' \simeq q''$ **then**
          $Q_{\text{REW}} := Q_{\text{REW}} \cup \{\langle q', \mathsf{r}, \mathsf{u}\rangle\}$;
        end
      end
      /* factorization step */
      foreach $S \subseteq \mathit{body}(q)$ that is factorizable w.r.t. σ **do**
        $q' := \gamma_S(q)$;
        if there is no $(q'', \star, \star) \in Q_{\text{REW}}$ such that $q' \simeq q''$ **then**
          $Q_{\text{REW}} := Q_{\text{REW}} \cup \{\langle q', \mathsf{f}, \mathsf{u}\rangle\}$;
        end
      end
    end
    /* query q is now explored */
    $Q_{\text{REW}} := (Q_{\text{REW}} \setminus \{(q, x, \mathsf{u})\}) \cup \{(q, x, \mathsf{e})\}$;
  end
until $Q_{\text{TEMP}} = Q_{\text{REW}}$;
$Q_{\text{FIN}} := \{q \mid \langle q, \mathsf{r}, \mathsf{e}\rangle \in Q_{\text{REW}},$ and q contains only predicates of $\mathbf{S}\}$;
return Q_{FIN}

The Algorithm XRewrite

In view of the fact that the rewriting algorithm XRewrite is heavily used in our complexity analysis, we would like to recall its definition. This algorithm is based on resolution, and thus, before we proceed further, we need to recall the crucial notion of unification. A set of atoms $A = \{\alpha_1, \ldots, \alpha_n\}$, where $n \geqslant 2$, *unifies* if there exists a substitution γ, called *unifier* for A, such that $\gamma(\alpha_1) = \cdots = \gamma(\alpha_n)$. A *most general unifier (MGU)* for A is a unifier for A, denoted as γ_A, such that for each other unifier γ for A, there exists a substitution γ' such that $\gamma = \gamma' \circ \gamma_A$. Notice that if a set of atoms unify, then there exists a MGU. Furthermore, the MGU for a set of atoms is unique (modulo variable renaming).

The algorihtm proceeds by exhaustively applying two steps: *rewriting* and *factorization*, which in turn rely on the technical notions of *applicability* and *factorizability*, respectively. We assume, w.l.o.g., that tgds and CQs do not share variables. Given a CQ q, a variable x is called *shared* in q if x is a free variable of q, or it occurs more than once in q. In what follows, we assume, w.l.o.g., that tgds are in normal form, i.e., they have only one atom in the head, and only one occurrence of an existentially quantified variable [27]. We write $\pi_\exists(\sigma)$ for the position at which the existentially quantified variable of σ occurs; in case σ does not mention an existentially quantified variable, then $\pi_\exists(\sigma) = \varepsilon$. (Recall that a position $P[i]$ identifies the i-th attribute of a predicate P.) We are now ready to recall applicability and factorizability; in what follows, we write $\mathit{body}(q)$ for the set of atoms occurring in q, and $\mathit{head}(\sigma)$ for the head-atom of σ.

Definition 6. (Applicability) Consider a CQ q and a tgd σ. Given a set of atoms $S \subseteq \mathit{body}(q)$, we say that σ is *applicable* to S if the following conditions are satisfied:

1. the set $S \cup \{\mathit{head}(\sigma)\}$ unifies, and
2. for each $\alpha \in S$, if the term at position π in α is either a constant or a shared variable in q, then $\pi \neq \pi_\exists(\sigma)$. ■

Roughly, whenever σ is *applicable* to S, this means that the atoms of S may be generated during the chase procedure by applying σ. Therefore, we are allowed to apply a rewriting step (which is essentially a resolution step) that resolves S using σ, i.e., S is replaced by $\mathit{body}(\sigma)$, and a new CQ that is closer to the input database is obtained.

If we start applying rewriting steps blindly, without checking for applicability, then the soundness of the rewriting procedure is not guaranteed. However, it is possible that the applicability condition is not satisfied, but still we should apply a rewriting step. This may happen due to the presence of redundant atoms in a query. For example, given the CQ

$$q \;=\; \exists x \exists y \exists z\, (R(x, y) \;\wedge\; R(x, z))$$

and the tgd

$$\sigma \;=\; P(u, v) \rightarrow \exists w\, R(w, u)$$

the applicability condition fails since the shared variable x in q occurs at the position $\pi_\exists(\sigma) = R[1]$. However, q is essentially the CQ $q = \exists x \exists y\, R(x, y)$, and now the applicability condition is satisfied. From the above informal discussion, we conclude that the applicability condition may prevent the algorithm from being complete since some valid rewriting steps are blocked. Because of this reason, we need the so-called factorization step, which aims at converting some shared variables into non-shared variables, and thus, satisfy the applicability condition. In general, this can be achieved by exhaustively unifying all the atoms that unify in the body of a CQ. However, some of these unifications do not contribute in any way to satisfying the applicability condition, and, as a result, many superfluous CQs are generated. It is thus better to apply a restricted form of factorization that generates a possibly small number of CQs that are vital for the completeness of the rewriting algorithm. This corresponds to the identification of all the atoms in the query whose shared existential variables come from the same atom in the chase, and they can be unified with no loss of information. Summing up, the key idea underlying the notion of factorizability is as follows: in order to apply the factorization step, there must exist a tgd that can be applied to its output.[7]

Definition 7. (Factorizability) Consider a CQ q and a tgd σ. Given a set of atoms $S \subseteq \mathit{body}(q)$, where $|S| \geqslant 2$, we say that S is *factorizable* w.r.t. σ if the following conditions are satisfied:

1. S unifies,
2. $\pi_\exists(\sigma) \neq \varepsilon$, and
3. there exists a variable $x \notin \mathit{var}(\mathit{body}(q) \setminus S)$ that occurs in every atom of S only at position $\pi_\exists(\sigma)$. ∎

Having the above key notions in place, we are now ready to recall the algorithm XRewrite, which is depicted in Algorithm 1. As said above, the UCQ rewriting of an OMQ $q = (\mathbf{S}, \Sigma, q)$ is computed by exhaustively applying (i.e., until a fixpoint is reached) the rewriting and the factorization steps. Notice that the CQs that are the result of the factorization step, are nothing else than auxiliary queries which are critical for the completeness of the final rewriting, but are not needed in the final rewriting. Thus, during the iterative procedure, the queries are labeled with r (resp., f) in order to keep track which of them are generated by the rewriting (resp., factorization) step. The CQ that is part of the input OMQ, although is not a result of the rewriting step, is labeled by r since it must be part of the final rewriting. Moreover, once the two crucial steps have been exhaustively applied on a CQ q, it is not necessary to revisit q since this will lead to redundant queries. Hence, the queries are also labeled with e (resp., u) indicating that a query has been already explored (resp., is unexplored). Let us now describe the two main steps of the algorithm. In the sequel, consider a triple (q, x, y), where $(x, y) \in \{\mathsf{r}, \mathsf{f}\} \times \{\mathsf{e}, \mathsf{u}\}$ (this is how we indicate that q is labeled by x and y), and a tgd $\sigma \in \Sigma$. We assume that q is of the form $\exists \bar{x}\, \varphi(\bar{x}, \bar{y})$.

Rewriting Step. For each $S \subseteq \mathit{body}(q)$ such that σ is applicable to S, the i-th application of the rewriting step generates the query $q' = \gamma_{S,\sigma^i}(q[S/\mathit{body}(\sigma^i)])$, where σ^i is the tgd obtained from σ by replacing each variable x with x^i, γ_{S,σ^i} is the MGU for the set $S \cup \{\mathit{head}(\sigma^i)\}$ (which is the identity on the variables that appear in the body but not in the head of σ^i), and $q[S/\mathit{body}(\sigma^i)]$ is obtained from q be replacing S with $\mathit{body}(\sigma^i)$. By considering σ^i (instead of σ) we basically rename, using the integer i, the variables of σ. This renaming step is needed in order to avoid undesirable clutters among the variables introduced during different applications of the rewriting step. Finally, if there is no $(q'', \mathsf{r}, \star) \in Q_{\text{REW}}$, i.e., an (explored or unexplored) query that is the result of the rewriting step, such that q' and q'' are the same (modulo bijective variable renaming), denoted $q' \simeq q''$, then $(q', \mathsf{r}, \mathsf{u})$ is added to Q_{REW}.

Factorization Step. For each $S \subseteq \mathit{body}(q)$ that is factorizable w.r.t. σ, the factorization step generates the query $q' = \gamma_S(q)$, where γ_S is the MGU for S. If there is no $(q'', \star, \star) \in Q_{\text{REW}}$, i.e., a query that is the result of the rewriting or the factorization step, and is explored or unexplored, such that $q' \simeq q''$, then $(q', \mathsf{f}, \mathsf{u})$ is added to Q_{REW}.

Proof of Proposition 14

We assume, w.l.o.g., that the predicates of $\mathbf{S}$ do not appear in the head of a tgd of Σ. Since $\Sigma \in \mathbb{NR}$, by Lemma 32, Σ admits a stratification $\{\Sigma_1, \ldots, \Sigma_n\}$ with stratification function $\mu : \mathit{sch}(\Sigma) \to \{0, \ldots, n\}$. Let us briefly explain how the rewriting tree T_Q of the OMQ $Q = (\mathbf{S}, \Sigma, q)$ is defined. T_Q is a rooted tree with q being its root. The i-th level of T_Q consists of the CQs obtained from the CQs of the $(i-1)$-th level by applying rewriting steps (see the algorithm XRewrite for details on the rewriting step) using only tgds from Σ_{n-i+1}. It is easy to verify that the CQs of the i-th level contain only predicates P such that $\mu(P) < n - i + 1$. It is now clear that the n-th level of T_Q (i.e., the leaves of T_Q) consists only of CQs obtained during the execution of XRewrite(Q) that contain only predicates of $\mathbf{S}$. Thus, in order to obtained the desired upper bound, it suffices to show that the number of atoms that occur in a CQ that is a leaf of T_Q is at most $|q| \cdot (\max_{\tau \in \Sigma}\{|\mathit{body}(\tau)|\})^{|\mathit{sch}(\Sigma)|}$. To this end, let us focus on one branch B of T_Q from the root q to a leaf q'. Such a branch can be naturally represented as a k-ary forest F_Q^B, where the root nodes are the atoms of q, and whenever an atom α is resolved during the rewriting step using a tgd τ, the atoms of $\mathit{body}(\tau)$, after applying the appropriate MGU, are the child nodes of α. Therefore, to obtain the desired upper bound, it suffices to show that the number of leaves of F_Q^B is at most $|q| \cdot (\max_{\tau \in \Sigma}\{|\mathit{body}(\tau)|\})^{|\mathit{sch}(\Sigma)|}$. By construction, F_Q^B consists, in general, of $|q|$ k-ary rooted trees, where $k = \max_{\tau \in \Sigma}\{|\mathit{body}(\tau)|\}$, of depth n. Hence, the number of leaves of F_Q^B is at most $|q| \cdot (\max_{\tau \in \Sigma}\{|\mathit{body}(\tau)|\})^n$. Since $n \leq |\mathit{sch}(\Sigma)|$, the claim follows.

[7] Let us clarify that for the purposes of the present work we can rely on the naive approach of exhaustively unifying all the atoms that unify in the body of a CQ. However, we would like to be consistent with [40], where the algorithm XRewrite is proposed, and thus, we stick on the slightly more involved notion of factorizability.

Proof of Theorem 16

A proof sketch for the coNEXPTIME$^{\text{NP}}$ upper bound is given in the main body of the paper. We proceed to establish the P$^{\text{NEXP}}$-hardness. Our proof is by reduction from a tiling problem that has been recently introduced in [34], which in turn relies on the standard *Exponential Tiling Problem*. Let us first recall the latter problem.

An instance of the Exponential Tiling Problem is a tuple (n, m, H, V, s), where n, m are numbers (in unary), H, V are subsets of $\{1, \ldots, m\} \times \{1, \ldots, m\}$, and s is a sequence of numbers of $\{1, \ldots, m\}$. Such a tuple specifies that we desire a $2^n \times 2^n$ grid, where each cell is tiled with a tile from $\{1, \ldots, m\}$. H (resp., V) is the horizontal (resp., vertical) compatibility relation, while s represents a constraint on the initial part of the first row of the grid. A *solution* to such an instance of the Exponential Tiling Problem is a function $f : \{0, \ldots, 2^n - 1\} \times \{0, \ldots, 2^n - 1\} \to \{1, \ldots, m\}$ such that:

1. $f(i, 0) = s[i]$, for each $0 \leq i \leq (|s| - 1)$;
2. $(f(i, j), f(i + 1, j)) \in H$, for each $0 \leq i \leq 2^n - 2$ and $0 \leq j \leq 2^n - 1$; and
3. $(f(i, j), f(i, j + 1)) \in V$, for each $0 \leq i \leq 2^n - 1$ and $0 \leq j \leq 2^n - 2$.

We will refer to $\{0, \ldots, 2^n - 1\} \times \{0, \ldots, 2^n - 1\}$ as a *grid*, with the pairs in it being *cells*. A cell consists of two *coordinates*, the column-coordinate (for short col-coordinate) and the row-coordinate, and any function on a grid is a *tiling*. The Exponential Tiling Problem is defined as follows: given an instance T as above, decide whether T has a solution. It is known that this problem is NEXPTIME-hard (see, e.g., Section 3.2 of [46]).

We are now ready to recall the tiling problem introduced in [34], called *Extended Tiling Problem* (ETP), which is P$^{\text{NEXP}}$-hard. An instance of this problem is a tuple $(k, n, m, H_1, V_1, H_2, V_2)$, where k, n, m are numbers (in unary), and H_1, V_1, H_2, V_2 are subsets of $\{1, \ldots, m\} \times \{1, \ldots, m\}$. The question is as follows: is it the case that for every sequence s, where $|s| = k$, of numbers of $\{1, \ldots, m\}$, (n, m, H_1, V_1, s) has no solution *or* (n, m, H_2, V_2, s) has a solution?

We give a reduction from the ETP to $\mathsf{Cont}(\mathbb{NR}, \mathbb{CQ})$. More precisely, given an instance $T = (k, n, m, H_1, V_1, H_2, V_2)$ of the ETP, our goal is to construct in polynomial time two queries $Q_i = (\mathbf{S}, \Sigma_i, q_i) \in (\mathbb{NR}, \mathbb{CQ})$, for $i \in \{1, 2\}$, such that T has a solution iff $Q_1 \subseteq Q_2$.

Data Schema S

The data schema $\mathbf{S}$ consists of:

- 0-ary predicates C_i^j, for each $i \in \{0, \ldots, k - 1\}$ and $j \in \{1, \ldots, m\}$; the atom C_i^j indicates that $s_i = j$.

The Query Q_1

The goal of the query Q_1 is twofold: (i) to check that the so-called *existence property* of the input database, i.e., for every $i \in \{0, \ldots, k - 1\}$, there exists at least one atom of the form C_i^j, is satisfied, and (ii) to check whether (n, m, H_1, V_1, s), where s is the sequence of tilings encoded in the input database, has a solution. To this end, the query Q_1 will mention the following predicates:

- 0-ary predicate C_i, indicating that there exists at least one atom of the form C_i^j in the input database.
- 0-ary predicate Existence, indicating that the input database enjoys the existence property.
- Unary predicate Tile_i, for each $i \in \{1, \ldots, m\}$; the atom $\text{Tile}_i(x)$ states that x is the tile i.
- Binary predicate H; the atom $H(x, y)$ encodes the fact that $(x, y) \in H_1$.
- Binary predicate V; the atom $V(x, y)$ encodes the fact that $(x, y) \in V_1$.
- 5-ary predicate T_i, for each $i \in \{1, \ldots, n\}$; the atom $T_i(x, x_1, x_2, x_3, x_4)$ states that x is a $2^i \times 2^i$ tiling obtained from the $2^{i-1} \times 2^{i-1}$ tilings $x_1, \ldots, x_4$ – details on the inductive construction of $2^i \times 2^i$ tilings from $2^{i-1} \times 2^{i-1}$ tilings are given below.
- Unary predicate Initial_i, for each $i \in \{0, \ldots, k - 1\}$; the atom $\text{Initial}_i(x)$ states that $s[i] = x$, i.e., the i-th element of the sequence s is x.
- Binary predicate Top_i^j, for each $i \in \{1, \ldots, n\}$ and $j \in \{0, \ldots, k - 1\}$; the atom $\text{Top}_i^j(x, y)$ states that in the $2^i \times 2^i$ tiling x the tile at position $(j, 0)$ is y.
- 0-ary predicate Tiling, indicating that there exists a $2^n \times 2^n$ tiling that is compatible with the initial tiling s encoded in the input database.
- 0-ary predicate Goal, which is derived whenever the predicates Existence and Tiling are derived.

Q_1 is defined as the query $(\mathbf{S}, \Sigma_1, \text{Goal})$, where Σ_1 consists of the following tgds:

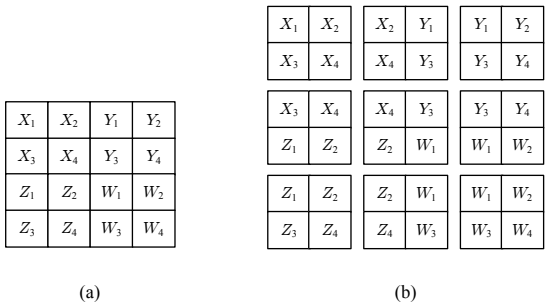

Figure 2: Inductive construction of tilings.

- Checking for the existence property of the input database

  For each $i \in \{0, \ldots, k-1\}$ and $j \in \{1, \ldots, m\}$:

  $$C_i^j \quad \rightarrow \quad C_i$$

  and the tgd that checks for the existence property

  $$C_0, \ldots, C_{k-1} \quad \rightarrow \quad \text{Existence}$$

- Generate the tiles

  $$\rightarrow \quad \exists x_1 \ldots \exists x_m \, (\text{Tile}_1(x_1), \ldots, \text{Tile}_m(x_m))$$

- Generate the compatibility relations

  For each $(i, j) \in H_1$:

  $$\text{Tile}_i(x), \text{Tile}_j(y) \quad \rightarrow \quad H(x, y)$$

  For each $(i, j) \in V_1$:

  $$\text{Tile}_i(x), \text{Tile}_j(y) \quad \rightarrow \quad V(x, y)$$

- Generate the $2^n \times 2^n$ tilings. The key idea is to inductively construct $2^i \times 2^i$ tilings from $2^{i-1} \times 2^{i-1}$ tilings. It is easy to verify that the grid in Figure 2(a) is a $2^i \times 2^i$ tiling iff the nine subgrids of it, shown in Figure 2(b), are $2^{i-1} \times 2^{i-1}$ tilings. This has been already observed in [33], where Datalog with complex values is studied.

  First, we construct tilings of size 2×2 (the base case of the inductive construction):

  $$H(x_1, x_2), H(x_3, x_4), V(x_1, x_3), V(x_2, x_4) \quad \rightarrow \quad \exists x \, T_1(x, x_1, x_2, x_3, x_4)$$

  Then, we inductively construct tilings of larger size until we get tilings of size $2^n \times 2^n$. This is done using the following tgds. For each $i \in \{2, \ldots, n\}$:

  $$\begin{aligned} &T_{i-1}(x_1, x_{11}, x_{12}, x_{21}, x_{22}), T_{i-1}(x_2, x_{12}, x_{13}, x_{22}, x_{23}), T_{i-1}(x_3, x_{13}, x_{14}, x_{23}, x_{24}) \\ &T_{i-1}(x_4, x_{21}, x_{22}, x_{31}, x_{32}), T_{i-1}(x_5, x_{22}, x_{23}, x_{32}, x_{33}), T_{i-1}(x_6, x_{23}, x_{24}, x_{33}, x_{34}), \\ &T_{i-1}(x_7, x_{31}, x_{32}, x_{41}, x_{42}), T_{i-1}(x_8, x_{32}, x_{33}, x_{42}, x_{43}), T_{i-1}(x_9, x_{33}, x_{34}, x_{43}, x_{44}) \rightarrow \\ &\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \exists x \, T_i(x, x_1, x_3, x_7, x_9) \end{aligned}$$

- Extract from the $2^n \times 2^n$ tilings the tiles at positions $(0, 0), (1, 0), \ldots, (k-1, 0)$. This is done using the following tgds:

  $$\begin{aligned} T_1(x, x_1, x_2, x_3, x_4) \quad &\rightarrow \quad \text{Top}_1^0(x, x_1), \text{Top}_1^1(x, x_2) \\ T_2(x, x_1, x_2, x_3, x_4), \text{Top}_1^0(x_1, y_0), \text{Top}_1^1(x_1, y_1) \quad &\rightarrow \quad \text{Top}_2^0(x, y_0), \text{Top}_2^1(x, y_1) \\ T_2(x, x_1, x_2, x_3, x_4), \text{Top}_1^0(x_2, y_0), \text{Top}_1^1(x_2, y_1) \quad &\rightarrow \quad \text{Top}_2^2(x, y_0), \text{Top}_2^3(x, y_1) \\ &\vdots \\ T_\ell(x, x_1, x_2, x_3, x_4), \text{Top}_{\ell-1}^0(x_1, y_0), \cdots, \text{Top}_{\ell-1}^{2^{\ell-1}-1}(x_1, y_{2^{\ell-1}-1}) \quad &\rightarrow \quad \text{Top}_\ell^0(x, y_0), \cdots, \text{Top}_\ell^{2^{\ell-1}-1}(x, y_{2^{\ell-1}-1}) \\ T_\ell(x, x_1, x_2, x_3, x_4), \text{Top}_{\ell-1}^0(x_1, y_0), \cdots, \text{Top}_{\ell-1}^{k-2^{\ell-1}-1}(x_1, y_{k-2^{\ell-1}-1}) \quad &\rightarrow \quad \text{Top}_\ell^{2^{\ell-1}}(x, y_0), \cdots, \text{Top}_\ell^{k-1}(x, y_{k-2^{\ell-1}-1}), \end{aligned}$$

  where $\ell = \lceil \log k \rceil$. Moreover, for each $i \in \{\ell + 1, \ldots, n\}$:

  $$T_i(x, x_1, x_2, x_3, x_4), \text{Top}_{i-1}^0(x_1, y_0), \cdots, \text{Top}_{i-1}^{k-1}(x_1, y_{k-1}) \quad \rightarrow \quad \text{Top}_i^0(x, y_0), \ldots, \text{Top}_i^{k-1}(x, y_{k-1})$$

- Check whether there exists a $2^n \times 2^n$ tiling that is compatible with the sequence of tilings s

  For each $i \in \{0, \ldots, k-1\}$ and $j \in \{1, \ldots, m\}$:

  $$C_i^j, \text{Tile}_j(x) \quad \rightarrow \quad \text{Initial}_i(x)$$

  and the tgd

  $$\text{Top}_n^0(x, y_0), \text{Initial}_0(y_0), \cdots, \text{Top}_n^{k-1}(x, y_{k-1}), \text{Initial}_{k-1}(y_{k-1}) \quad \rightarrow \quad \text{Tiling}$$

- Finally, we have the output tgd

  $$\text{Existence}, \text{Tiling} \quad \rightarrow \quad \text{Goal}$$

This concludes the construction of Q_1.

The Query Q_2

The goal of the query Q_2 is twofold: (i) to check that the so-called *uniqueness property* of the input database, i.e., for every $i \in \{0, \ldots, k-1\}$, there exists at most one atom of the form C_i^j, is satisfied, and (ii) to check whether (n, m, H_2, V_2, s), where s is the sequence of tilings encoded in the input database, has a solution. The query Q_2 mentions the same predicates as Q_1, and is defined as $(\mathbf{S}, \Sigma_2, \text{Goal})$, where Σ_2 consists of the following tgds:

- Checking the uniqueness property

  For each $i \in \{0, \ldots, k-1\}$ and $j, \ell \in \{1, \ldots, m\}$ with $j < \ell$:

  $$C_i^j, C_i^\ell \quad \rightarrow \quad \text{Goal}$$

- The rest of Σ_2 encodes the tiling problem (n, m, H_2, V_2, s) in exactly the same way as Σ_1 encodes (n, m, H_1, V_1, s).

This concludes the construction of Q_2.

Proof of Proposition 18

The set Σ^n consists of the following tgds; for brevity, we write $\bar{x}_i^j$ for $x_i, x_{i+1}, \ldots, x_j$:[8]

$$\begin{aligned} S(x_1, \ldots, x_n) \quad &\rightarrow \quad P_n(x_1, \ldots, x_n) \\ P_i(\bar{x}_1^{i-1}, z, \bar{x}_{i+1}^n, z, o), P_i(\bar{x}_1^{i-1}, o, \bar{x}_{i+1}^n, z, o) \quad &\rightarrow \quad P_{i-1}(\bar{x}_1^{i-1}, z, \bar{x}_{i+1}^n, z, o), \quad 1 \leq i \leq n, \\ P_0(\underbrace{z, \ldots, z}_{n}, z, o) \quad &\rightarrow \quad \text{Ans}(z, o), \end{aligned}$$

while $q = \text{Ans}(0, 1)$. It can be verified that, for every $\{S\}$-database D, $Q^n(D) \neq \varnothing$ implies that

$$D \supseteq \{S(c_1, \ldots, c_{n-2}, 0, 1) \mid (c_1, \ldots, c_{n-2}) \in \{0, 1\}^{n-2}\},$$

and thus, $|D| \geq 2^{n-2}$. Let $Q = (\{S\}, \Sigma', q')$, where Σ' is a set of tgds and q' a Boolean CQ, and D an $\{S\}$-database. Clearly, $Q^n(D) \not\subseteq Q(D)$ iff $Q^n(D) \neq \varnothing$ and $Q(D) = \varnothing$. This implies that $|D| \geq 2^{n-2}$, and the claim follows.

Proof of Theorem 19

The coNExpTime upper bound, as well as the Π_2^P-hardness in case of fixed-arity predicates, are discussed in the main body of the paper. Here, we show the coNExpTime-hardness. The proof proceeds in two steps:

1. First, we show that $\mathsf{Cont}((\mathbb{FNR}, \mathbb{CQ}), (\mathbb{L}, \mathbb{UCQ}))$ is coNExpTime-hard, where $\mathbb{FNR}$ denotes the class of full non-recursive sets of tgds, i.e., non-recursive sets of tgds without existentially quantified variables.
2. Then, we reduce $\mathsf{Cont}((\mathbb{FNR}, \mathbb{CQ}), (\mathbb{L}, \mathbb{UCQ}))$ to $\mathsf{Cont}((\mathbb{S}, \mathbb{CQ}), (\mathbb{L}, \mathbb{UCQ}))$ by showing that (under some assumptions that are explained below) every query in $(\mathbb{FNR}, \mathbb{CQ})$ can be rewritten as an $(\mathbb{S}, \mathbb{CQ})$ query.

By Proposition 9, we immediately get that $\mathsf{Cont}((\mathbb{S}, \mathbb{CQ}), (\mathbb{L}, \mathbb{CQ}))$ is coNExpTime-hard, as needed.

Step 1: $\mathsf{Cont}((\mathbb{FNR}, \mathbb{CQ}), (\mathbb{L}, \mathbb{UCQ}))$ is coNExpTime-hard

We show that $\mathsf{Cont}((\mathbb{FNR}, \mathbb{CQ}), (\mathbb{L}, \mathbb{UCQ}))$ is coNExpTime-hard, even if we focus on *0-1 queries*, that is, queries Q with following property: for every database D, $Q(D) = Q(D_{01})$, where $D_{01} \subseteq D$ is the restriction of D on the binary domain $\{0, 1\}$, i.e., $D_{01} = \{R(\bar{c}) \in D \mid \bar{c} \subseteq \{0, 1\}\}$. The proof is by reduction from the Exponential Tiling Problem, and is a non-trivial adaptation of the one given in [14] for showing that containment of non-recursive Datalog queries is coNExpTime-hard.

Theorem 34. $\mathsf{Cont}((\mathbb{FNR}, \mathbb{CQ}), (\mathbb{L}, \mathbb{UCQ}))$ *is co*NExpTime*-hard, even for 0-1 queries.*

Proof. Given an instance $T = (n, m, H, V, s)$ of the Exponential Tiling Problem, we are going to construct a $(\mathbb{FNR}, \mathbb{CQ})$ 0-1 query $Q_T = (\mathbf{S}, \Sigma, q)$ and a $(\mathbb{L}, \mathbb{UCQ})$ 0-1 query $Q_T' = (\mathbf{S}, \Sigma_T, q_T)$ such that T has a solution iff $Q_T \not\subseteq Q_T'$.

[8] A similar construction has been used in [40] for showing a lower bound on the size of a CQ in the UCQ rewriting of a $(\mathbb{S}, \mathbb{CQ})$ OMQ.

Data Schema **S**

The data schema **S** consists of:

- $2n$-ary predicates TiledBy_i, for each $i \leq m$; the atom $\text{TiledBy}_i(x_1, \ldots, x_n, y_1, \ldots, y_n)$ indicates that the cell with coordinates $((x_1, \ldots, x_n), (y_1, \ldots, y_n)) \in \{0,1\}^n \times \{0,1\}^n$ is tiled by tile i. Notice that we use n-bit binary numbers to represent a coordinate; this is the key difference between our construction and the one of [14].

The Query Q_T

The goal of the query Q_T is to assert whether the input database encodes a candidate tiling, i.e., whether the entire grid is tiled, without taking into account the constraints, that is, the compatibility relations and the constraint on the initial part of the first row. To this end, the query Q_T will mention the following predicates:

- Unary predicate Bit; the atom $\text{Bit}(x)$ simply says that x is a bit, i.e., $x \in \{0,1\}$.
- $2n$-ary predicate TiledAboveCol_i, for each $i \leq n$; the atom $\text{TiledAboveCol}_i(\bar{x}, \bar{y})$ says that for the row-coordinate $\bar{y}$ there are tiled cells with coordinates $(\bar{x}', \bar{y})$ for every col-coordinate $\bar{x}'$ that agrees with $\bar{x}$ on the first $i-1$ bits. In other words, for the row corresponding to $\bar{y}$, every column extending the first $i-1$ bits of $\bar{x}$ is tiled. In particular, $\text{TiledAboveCol}_1(\bar{x}, \bar{y})$ says that the entire row $\bar{y}$ is tiled.
- $2n$-ary predicate TiledAboveRow_i, for each $i \leq n$; the atom $\text{TiledAboveRow}_i(\bar{y})$ says that for every $\bar{y}'$ that agrees with $\bar{y}$ on the first $i-1$ bits, the row $\bar{y}'$ is fully tiled.
- n-ary predicate RowTiled; the atom $\text{RowTiled}(\bar{y})$ says that the row $\bar{y}$ is fully tiled.
- 0-ary predicate AllTiled, which asserts that the entire grid is tiled.
- 0-ary predicate Goal, which is derived whenever the predicate AllTiled is derived.

Q_T is defined as the query $(\mathbf{S}, \Sigma, \text{Goal})$, where Σ consists of the following rules:

- Generate Bit atoms

$$\begin{aligned} &\rightarrow \quad \text{Bit}(0) \\ &\rightarrow \quad \text{Bit}(1). \end{aligned}$$

- RowTiled

  For each $j, k \leq m$:

$$\text{TiledBy}_j(x_1, \ldots, x_{n-1}, 1, y_1, \ldots, y_n), \text{TiledBy}_k(x_1, \ldots, x_{n-1}, 0, y_1, \ldots, y_n), \\ \text{Bit}(x_1), \ldots, \text{Bit}(x_{n-1}), \text{Bit}(y_1), \ldots, \text{Bit}(y_n), \text{Bit}(w) \rightarrow \\ \text{TiledAboveCol}_n(x_1, \ldots, x_{n-1}, w, y_1, \ldots, y_n)$$

  For each $2 \leq i \leq n$:

$$\text{TiledAboveCol}_i(x_1, \ldots, x_{i-1}, 1, x_{i+1}, \ldots, x_n, y_1, \ldots, y_n), \\ \text{TiledAboveCol}_i(x_1, \ldots, x_{i-1}, 0, x'_{i+1}, \ldots, x'_n, y_1, \ldots, y_n), \\ \text{Bit}(w_i), \ldots, \text{Bit}(w_n) \rightarrow \\ \text{TiledAboveCol}_{i-1}(x_1, \ldots, x_{i-1}, w_i, \ldots, w_n, y_1, \ldots, y_n)$$

  A row is fully tiled:

$$\text{TiledAboveCol}_1(x_1, \ldots, x_n, y_1, \ldots, y_n) \quad \rightarrow \quad \text{RowTiled}(y_1, \ldots, y_n)$$

- AllTiled

$$\text{RowTiled}(y_1, \ldots, y_{n-1}, 1), \text{RowTiled}(y_1, \ldots, y_{n-1}, 0), \text{Bit}(w) \quad \rightarrow \quad \text{TiledAboveRow}_n(y_1, \ldots, y_{n-1}, w)$$

  For each $2 \leq i \leq n$:

$$\text{TiledAboveRow}_i(y_1, \ldots, y_{i-1}, 1, y_{i+1}, \ldots, y_n), \\ \text{TiledAboveRow}_i(y_1, \ldots, y_{i-1}, 0, y'_{i+1}, \ldots, y'_n), \\ \text{Bit}(w_i), \ldots, \text{Bit}(w_n) \rightarrow \\ \text{TiledAboveRow}_{i-1}(y_1, \ldots, y_{i-1}, w_i, \ldots, w_n)$$

  The entire grid is tiled:

$$\begin{aligned} \text{TiledAboveRow}_1(y_1, \ldots, y_n) \quad &\rightarrow \quad \text{AllTiled} \\ \text{AllTiled} \quad &\rightarrow \quad \text{Goal} \end{aligned}$$

This concludes the construction of the query Q_T.

The Query Q'_T

Q'_T is defined in such a way that $Q'_T(D)$ is non-empty exactly when the input database D encodes an invalid tiling, i.e., when one of the constraints on the tiles is violated. The query Q'_T will mention the following intensional predicates:

- Unary predicate Bit; as above, $\text{Bit}(x)$ says that x is a bit.
- $2i$-ary predicate LastFirst_i, for each $1 \leq i \leq n$; the atom $\text{LastFirst}_i(x_1, \ldots, x_i, y_1, \ldots, y_i)$ says that $(x_1, \ldots, x_i) = (1, \ldots, 1)$ and $(y_1, \ldots, y_i) = (0, \ldots, 0)$.
- $2i$-ary predicate Succ_i, for each $1 \leq i \leq n$; the atom $\text{Succ}_i(\bar{x}, \bar{y})$ says that the i-bit binary number $\bar{y}$ is the successor of the i-bit binary number $\bar{x}$.
- 0-ary predicate Goal.

Q'_T is defined as the query $(\mathbf{S}, \Sigma', q')$. The set Σ' consists of the following linear tgds:

- Generate Bit atoms:

$$\begin{aligned} &\rightarrow \text{Bit}(0) \\ &\rightarrow \text{Bit}(1). \end{aligned}$$

- Generate the successor predicates:

$$\begin{aligned} &\rightarrow \text{Succ}_1(0, 1) \\ &\rightarrow \text{LastFirst}_1(1, 0). \end{aligned}$$

  For each $1 \leq i \leq n-1$:

$$\begin{aligned} \text{Succ}_i(x_1, \ldots, x_i, y_1, \ldots, y_i) &\rightarrow \text{Succ}_{i+1}(0, x_1, \ldots, x_i, 0, y_1, \ldots, y_i) \\ \text{Succ}_i(x_1, \ldots, x_i, y_1, \ldots, y_i) &\rightarrow \text{Succ}_{i+1}(1, x_1, \ldots, x_i, 1, y_1, \ldots, y_i) \\ \text{LastFirst}_i(x_1, \ldots, x_i, y_1, \ldots, y_i) &\rightarrow \text{Succ}_{i+1}(0, x_1, \ldots, x_i, 1, y_1, \ldots, y_i) \\ \text{LastFirst}_i(x_1, \ldots, x_i, y_1, \ldots, y_i) &\rightarrow \text{LastFirst}_{i+1}(1, x_1, \ldots, x_i, 0, y_1, \ldots, y_i). \end{aligned}$$

The UCQ q' consists of the following (Boolean) CQs; for brevity, the existential quantifiers in front of the CQs are omitted:

- Tile Consistency

  For each $i \neq j \leq m$:

$$\text{TiledBy}_i(x_1, \ldots, x_n, y_1, \ldots, y_n), \text{TiledBy}_j(x_1, \ldots, x_n, y_1, \ldots, y_n), \\ \text{Bit}(x_1), \ldots, \text{Bit}(x_n), \text{Bit}(y_1), \ldots, \text{Bit}(y_n)$$

- Tile Compatibility

  For each $(i, j) \notin V$:

$$\text{Succ}_n(x_1, \ldots, x_n, y_1, \ldots, y_n), \\ \text{TiledBy}_i(w_1, \ldots, w_n, x_1, \ldots, x_n), \text{TiledBy}_i(w_1, \ldots, w_n, y_1, \ldots, y_n), \\ \text{Bit}(w_1), \ldots, \text{Bit}(w_n)$$

  For each $(i, j) \notin H$:

$$\text{Succ}_n(x_1, \ldots, x_n, y_1, \ldots, y_n), \\ \text{TiledBy}_i(x_1, \ldots, x_n, w_1, \ldots, w_n), \text{TiledBy}_i(y_1, \ldots, y_n, w_1, \ldots, w_n), \\ \text{Bit}(w_1), \ldots, \text{Bit}(w_n)$$

- Tiling of First Row

  For each $j \leq n$, let f_j be the function from $\{1, \ldots, n\}$ into $\{0, 1\}$ such that $f_j(1) \ldots f_j(n)$ is the number j in binary representation, and let $k \in \{1, \ldots, m\}$ other than $s[j]$; recall that s is a sequence of numbers of $\{1, \ldots, m\}$ that represents a constraint on the initial part of the first row of the grid. Then, we have the CQ:

$$\text{TiledBy}_k(x_1, \ldots, x_n, \underbrace{z, \ldots, z}_{n}), \text{Succ}_1(z, o)$$

  where, for each $i \in \{1, \ldots, n\}$, $x_i = z$ if $f_j(i) = 0$, and $x_i = o$ if $f_j(i) = 1$.

This concludes the definition of the query Q'_T. □

Step 2: $\mathsf{Cont}((\mathbb{S}, \mathbb{CQ}), (\mathbb{L}, \mathbb{UCQ}))$ is coNEXPTIME-hard

Our goal is show that every 0-1 query $(\mathbf{S}, \Sigma, q) \in (\mathbb{F}, \mathbb{CQ})$ can be equivalently rewritten as a 0-1 query $(\mathbf{S}, \Sigma', q')$, where all the tgds of Σ' are *lossless*, i.e., all the body-variables appear also in the head, which in turn implies that Σ' is sticky.

PROPOSITION 35. *Consider a 0-1 query $Q \in (\mathbb{F}, \mathbb{CQ})$. We can construct in polynomial time a 0-1 query $Q' \in (\mathbb{S}, \mathbb{CQ})$ such that $Q \equiv Q'$.*

PROOF. Let $Q = (\mathbf{S}, \Sigma, q)$, and assume that n is the maximum number of variables occurring in the body of a tgd of Σ. We are going to construct in polynomial time a 0-1 query $Q' = (\mathbf{S}, \Sigma', q') \in (\mathbb{S}, \mathbb{CQ})$ such that $Q \equiv Q'$.

The set Σ' consists of the following tgds:

- Initialization Rules

  We first transform every database atom of the form $R(\bar{c})$ into an atom $R'(\bar{c}, \underbrace{0, \ldots, 0}_{n}, 0, 1)$. This is done as follows:

  $$\begin{aligned} &\rightarrow \quad \mathrm{Bit}(0) \\ &\rightarrow \quad \mathrm{Bit}(1) \end{aligned}$$

  and, for each k-ary predicate $R \in \mathbf{S}$, we have the lossless tgd

  $$R(x_1, \ldots, x_k), \mathrm{Bit}(x_1), \ldots, \mathrm{Bit}(x_k) \quad \rightarrow \quad R'(x_1, \ldots, x_k, \underbrace{0, \ldots, 0}_{n})$$

  Notice that we can safely force the variables $x_1, \ldots, x_k$ to take only values from $\{0, 1\}$ due to the 0-1 property.

- Transformation into Lossless Tgds

  For each tgd $\sigma \in \Sigma$ of the form

  $$R_1(\bar{x}_1), \ldots, R_k(\bar{x}_k) \quad \rightarrow \quad R_0(\bar{x}_0)$$

  we have the lossless tgd

  $$R'_1(\bar{x}_1, \underbrace{0, \ldots, 0}_{n}), \ldots, R'_k(\bar{x}_k, \underbrace{0, \ldots, 0}_{n}) \quad \rightarrow \quad R'_0(\bar{x}_0, y_1, \ldots, y_n),$$

  where, if $\{v_1, \ldots, v_\ell\}$, for $\ell \in \{1, \ldots, n\}$, is the set of variables occurring in the body of σ (the order is not relevant), then $y_i = v_i$, for each $i \in \{1, \ldots, \ell\}$, and $y_j = v_1$, for each $j \in \{\ell + 1, \ldots, n\}$.

- Finalization Rules

  Observe that each atom obtained during the chase due to one of the lossless tgds introduced above is of the form $R'(\bar{x}, \bar{y})$, where $\bar{y} \in \{0, 1\}^n$. If $\bar{y} \neq (0, \ldots, 0)$, then we need to ensure that eventually the atom

  $$R'(\bar{x}, \underbrace{0, \ldots, 0}_{n})$$

  will be inferred. This is achieved by adding to Σ' the following tgds: For each k-ary predicate R occurring in Σ, and for each $1 \leq i \leq n$, we have the rule:

  $$R'(x_1, \ldots, x_k, y_1, \ldots, y_{i-1}, 1, y_{i+1}, \ldots, y_n) \quad \rightarrow \quad R'(x_1, \ldots, x_k, y_1, \ldots, y_{i-1}, 0, y_{i+1}, \ldots, y_n).$$

This concludes the definition of Σ'.

The CQ q' is defined analogously. More precisely, assuming that q is of the form (the existential quantifiers are omitted)

$$R_1(\bar{x}_1), \ldots, R_k(\bar{x}_k)$$

the CQ q' is defined as

$$R'_1(\bar{x}_1, \underbrace{0, \ldots, 0}_{n}), \ldots, R'_k(\bar{x}_k, \underbrace{0, \ldots, 0}_{n}).$$

It is easy to verify that Σ' consists of lossless tgds, and thus, $Q' \in (\mathbb{S}, \mathbb{CQ})$. It also not difficult to see that, for every database D over $\mathbf{S}$, $Q(D_{01}) = Q'(D_{01})$; thus, by the 0-1 property, $Q(D) = Q'(D)$, and the claim follows. □

By Theorem 34 and Proposition 35, we immediately get that $\mathsf{Cont}((\mathbb{S}, \mathbb{CQ}), (\mathbb{L}, \mathbb{UCQ}))$ is coNEXPTIME-hard, as needed.

PROOFS OF SECTION 5

Recall that, for the sake of technical clarity, we focus on constant-free tgds and CQs, but all the results can be extended to the general case at the price of more involved definitions and proofs. Moreover, we assume that tgds have only one atom in the head. This does not affect the generality of our proof since every set of guarded tgds can be transformed in polynomial time into a set of guarded tgds with the above property; see, e.g., [24]. Finally, for convenience of presentation, we also assume that the body of a tgd is non-empty, i.e., the body of a tgd is always an atom and not the symbol $\top$.

Proof of Proposition 21

Let us start by recalling the key notion of tree decomposition. Notice that the definition of the tree decomposition that we give here is slightly different than the one in the main body of the paper. The reason is because, for convenience of presentation, we prefer to employ a slightly different notation.

Definition 8. Let I be an instance. A *tree decomposition of I that omits V*, where $V \subseteq dom(I)$, is a pair $\delta = (\mathcal{T}, (X_t)_{t \in T})$, where $\mathcal{T} = (T, E^{\mathcal{T}})$ is a tree and $(X_t)_{t \in T}$ a family of subsets of $dom(I)$ (called the *bags* of the decomposition) such that:

1. For every $v \in dom(I) \setminus V$, the set $\{t \in T \mid v \in X_t\}$ is non-empty and connected.
2. For every atom $P(s_1, \ldots, s_n) \in I$, there is a $t \in T$ such that $\{s_1, \ldots, s_n\} \subseteq X_t$.

The *width* of a tree decomposition $\delta = (\mathcal{T}, (X_t)_{t \in T})$ omitting V is $\max\{|X_t| : t \in T\} - 1$. The *tree-width* of I is the minimum among the widths of all tree decompositions of I that omit V. We call a tree decomposition omitting $\varnothing$ simply *tree decomposition* of I. For $v \in T$, we denote by $I_\delta(v)$ the subinstance of I induced by X_v. ∎

Notation. We usually denote the strict partial order among the nodes of a tree $\mathcal{T}$ of a tree decomposition $\delta = (\mathcal{T}, (X_t)_{t \in T})$ by $\prec$. Accordingly, we write $v \preceq w$ iff $v \prec w$ or $v = w$. For brevity, ε will usually denote the root of a tree decomposition at hand. If ambiguities could possibly arise, we shall use subscripts in these notations. Furthermore, when δ is clear from context, we shall omit it from the expression $I_\delta(v)$.

Let $\delta = (\mathcal{T}, (X_t)_{t \in T})$ be a tree decomposition of I and $V \subseteq T$. Recall that δ is $[V]$-*guarded* (or *guarded except for V*), if for every node $v \in T \setminus V$, there is an atom $P(s_1, \ldots, s_n) \in I$ such that $X_v \subseteq \{s_1, \ldots, s_n\}$. A $[\varnothing]$-guarded tree decomposition of I is simply called *guarded tree decomposition*.

Also recall the crucial notion of C-tree:

Definition 9. An $\mathbf{S}$-instance I is a *C-tree*, where $C \subseteq I$, if there is a tree decomposition $\delta = (\mathcal{T}, (X_t)_{t \in T})$ of I such that

1. $I_\delta(\varepsilon) = C$, i.e., the subinstance of I induced by X_ε equals C.
2. δ is guarded except for $\{\varepsilon\}$.

If δ or C is clear from context, we shall often refer to $|dom(C)|$ as the *diameter* of D and to C as the *core* of D. ∎

Remark. The notion of C-tree defined here refers to both instances and databases, i.e., a C-tree may be a (finite) database or an instance. We often do not explicitly mention whether a C-tree at hand is a database or an instance. However, it will be clear from context whether a C-tree is a database or an instance.

We proceed to establish the following technical lemma, which in turn allows us to show Proposition 21. It is an adaption of a result in [7] to the case of guarded tgds. Henceforth, for brevity, given a query $Q = (\mathbf{S}, \Sigma, q) \in (\mathbb{G}, \mathbb{BCQ})$ and an $\mathbf{S}$-database D, we write $D \models Q$ for the fact that $Q(D) \neq \varnothing$.

LEMMA 36. *Let $Q = (\mathbf{S}, \Sigma, q)$ be an OMQ from $(\mathbb{G}, \mathbb{BCQ})$. Let D be an $\mathbf{S}$-database and suppose $D \models Q$. Then there is a finite $\mathbf{S}$-instance $\hat{I}$ such that $\hat{I} \models Q$ and:*

1. *$\hat{I}$ is a C-tree such that $|dom(C)| \leq ar(\mathbf{S} \cup sch(\Sigma)) \cdot |q|$.*
2. *There is a homomorphism from $\hat{I}$ to D.*

Before we proceed with its formal proof, let us explain why Proposition 21 is an easy consequence of Lemma 36. The fact that the first item implies the second is trivial. Conversely, suppose that $Q_1 \not\subseteq Q_2$, which implies that there exists an $\mathbf{S}$-database D such that $D \models Q_1$ and $D \not\models Q_2$. By Lemma 36, there exists a C-tree $\hat{I}$, where $|dom(C)| \leq ar(\mathbf{S} \cup sch(\Sigma_1)) \cdot |q_1|$, such that $\hat{I} \models Q_1$. Moreover, there is a homomorphism from $\hat{I}$ to D; hence, since Q_2 is closed under homomorphisms, it immediately follows that $\hat{I} \not\models Q_2$. Consequently, the $\mathbf{S}$-database $\hat{D}$ obtained from $\hat{I}$ after replacing each null z with a distinct constant c_z is a C-tree such that $Q_1(\hat{D}) \not\subseteq Q_2(\hat{D})$, and Proposition 21 follows.

We now proceed with the proof of Lemma 36 which is our main task in this section. Before that, we introduce some additional auxiliary concepts.

The Guarded Chase Forest

Given a database D and a set Σ of guarded tgds, the *guarded chase forest for D and Σ* is a forest (whose edges and nodes are labeled) constructed as follows:

1. For each fact $R(\bar{a})$ in D, add a node labeled with $R(\bar{a})$.
2. For each node v labeled with $\alpha \in \mathit{chase}(D, \Sigma)$ and for every atom β resulting from a one-step application of a rule $\tau \in \Sigma$, if α is the image of the guard in this application of τ, then add a node w labeled with β and introduce an arc from v to w labeled with τ.

We can assume that the guarded chase forest is always built inductively according to a fixed, deterministic version of the chase procedure. The non-root nodes are then totally ordered by a relation $\prec$ that reflects their order of generation. Furthermore, we can extend $\prec$ to database atoms by picking a lexicographic order among them. Notice that one atom can be the label of multiple nodes. Using the order $\prec$ we can, however, always refer to the $\prec$-least node.

Guarded Unraveling

Let I be an instance over $\mathbf{S}$. We say that $X \subseteq \mathit{dom}(I)$ is *guarded in I*, if there are $a_1, \ldots, a_s \in \mathit{dom}(I)$ such that

- $X \subseteq \{a_1, \ldots, a_s\}$ and
- there is an $R/s \in \mathbf{S}$ such that $I \models R(a_1, \ldots, a_s)$.

A tuple $\bar{t}$ is guarded in I if the set containing the elements of $\bar{t}$ is guarded in I.

In the following paragraph, we largely follow the notions introduced in [2, 13]. Fix an $\mathbf{S}$-instance I and some $X_0 \subseteq \mathit{dom}(I)$. Let Π be the set of finite sequences of the form $X_0 X_1 \cdots X_n$, where, for $i > 0$, X_i is a guarded set in I, and, for $i \geq 0$, $X_{i+1} = X_i \cup \{a\}$ for some $a \in \mathit{dom}(I) \setminus X_i$, or $X_i \supseteq X_{i+1}$. The sequences from Π can be arranged in a tree by their natural prefix order and each sequence $\pi = X_0 X_1 \cdots X_n$ identifies a unique node in this tree. In this context, we say that $a \in \mathit{dom}(I)$ is *represented at π* whenever $a \in X_n$. Two sequences π, π' are *a-equivalent*, if a is represented at each node on the unique shortest path between π and π'. For a represented at π, we denote by $[\pi]_a$ the a-equivalence class of π. The *guarded unraveling around X_0* is the instance I^* over the elements $\{[\pi]_a \mid a \text{ is represented at } \pi\}$, where

$$I^* \models R([\pi_1]_{a_1}, \ldots, [\pi_n]_{a_n}) \iff_{df} I \models R(a_1, \ldots, a_n) \text{ and } \exists \pi \in \Pi, \forall i \in \{1, \ldots, n\} \colon [\pi]_{a_i} = [\pi_i]_{a_i},$$

for all $R/n \in \mathbf{S}$.

LEMMA 37. *For every $\mathbf{S}$-instance I and any $X_0 \subseteq \mathit{dom}(I)$, the guarded unraveling I^* around X_0 is a C-tree over $\mathbf{S}$, where C is the subinstance of I^* induced by the elements $\{[X_0]_a \mid a \in X_0\}$.*

PROOF. Let $\delta = (\mathcal{T}, (X_t)_{t \in T})$, where $\mathcal{T}$ is the natural tree that arises from ordering the sequences in Π by their prefixes. For $\pi \in T$, let $X_\pi \coloneqq \{[\pi]_a \mid a \text{ is represented at } \pi\}$. Let ε denote the root of $\mathcal{T}$. We need to show that δ is an appropriate tree decomposition witnessing that I^* is a C-tree. First, note that it is clear that $I(\varepsilon) = \{[X_0]_a \mid a \in X_0\}$ by construction. Let $[\pi]_a \in \mathit{dom}(I)$ and consider the set $A \coloneqq \{t \in T \mid [\pi]_a \in X_t\}$. This set is certainly non-empty. Moreover, for $t_1, t_2 \in A$, we know that $[t_1]_a = [t_2]_a$, hence t_1 and t_2 are a-connected in $\mathcal{T}$. Suppose $I^* \models R([\pi_1]_{a_1}, \ldots, [\pi_n]_{a_n})$ for some $R/n \in \mathbf{S}$. Then there is a $\pi \in T$ such that $[\pi]_{a_i} = [\pi_i]_{a_i}$, for all $i = 1, \ldots, n$. Hence, $a_1, \ldots, a_n$ are all represented at π and so $\{[\pi_1]_{a_1}, \ldots, [\pi_n]_{a_n}\} \subseteq X_\pi$. It remains to show that δ is guarded except for $\{\varepsilon\}$. Let $\pi \neq \varepsilon$ and consider the set X_t. Since π is a sequence of length greater than one, its last element Y is a guarded set in I. Hence, there are $a_1, \ldots, a_s$ such that $Y \subseteq \{a_1, \ldots, a_s\}$ and $I \models R(a_1, \ldots, a_s)$ for some $R/s \in \mathbf{S}$. Let $\{a_1, \ldots, a_s\} \setminus Y = \{b_1, \ldots, b_m\}$ and define $\rho \coloneqq \pi \cdot (Y \cup \{b_1\}) \cdot (Y \cup \{b_1, b_2\}) \cdots (Y \cup \{b_1, \ldots, b_m\})$. Then $I^* \models R([\rho]_{a_1}, \ldots, [\rho]_{a_s})$, as desired. □

Notice that this lemma implies that the tree-width of I^* is bounded by $|X_0| + \mathit{ar}(\mathbf{S}) - 1$.

We are now ready to prove Lemma 36:

PROOF OF LEMMA 36. Let $q = \exists \bar{y}\, \varphi(\bar{y})$ and μ a homomorphism mapping $\varphi(\bar{y})$ to $\mathit{chase}(D, \Sigma)$. Let $R_1(\bar{b}_1), \ldots, R_k(\bar{b}_k)$ exhaust all facts from D that are the roots of those $\prec$-least facts from $\mu(\varphi(\bar{y}))$ in the guarded chase forest of D and Σ that have an element from $\mathit{dom}(D)$ as argument. Let $G_\mu \coloneqq \bigcup_{1 \leq i \leq k} \{\bar{b}_i\}$ and let I^* be the unraveling of D around G_μ, regarding all elements from $\mathit{dom}(I^*)$ as labeled nulls. Henceforth, for every $a \in G_\mu$, we denote by λ_a the element $[G_\mu]_a$. We say that λ_a *represents* a. Let C be the substructure of I^* induced by the set $\{\lambda_a \mid a \in G_\mu\}$. Notice that I^* is an infinite instance that is a C-tree by Lemma 37. We will show later how to get a finite instance from I^* that satisfies our constraints. We proceed to show that $I^* \cup \Sigma$ logically entails q, denoted $I^*, \Sigma \models q$:

LEMMA 38. $I^*, \Sigma \models q$.

PROOF. We will first construct a universal model J of I^* and Σ. Recall that an instance U is a *universal model* of I and Σ, if it can be homomorphically mapped to every model of $I \cup \Sigma$; in particular, it is well-known and easy to prove that $\mathit{chase}(I, \Sigma)$ is always a universal model of I and Σ. Before constructing J, we introduce some additional notions. In the following, given a guarded set $G = \{a_1, \ldots, a_k\}$ in D, a *copy of G in I^** is a set $\Gamma = \{\alpha_1, \ldots, \alpha_k\}$ which is guarded in I^* such

that, for $i = 1, \dots, k$, we have that $\alpha_i = [\pi_i]_{a_i}$ for some sequences π_i and $D \models R(a_{i_1}, \dots, a_{i_m})$ iff $I^* \models R(\alpha_{i_1}, \dots, \alpha_{i_m})$ for all $R \in \mathbf{S}$ and $i_j \in \{1, \dots, k\}$. Copies of guarded tuples are defined accordingly. Consider the structure $\mathit{chase}(D, \Sigma)$. Let G be a guarded set in D and $D \restriction G$ denote the subinstance of D induced by G. It is well-known and easy to prove that $\mathit{chase}(D \restriction G, \Sigma)$ is acyclic (cf., e.g., [23]). Henceforth, we loosely call $\mathit{chase}(D \restriction G, \Sigma)$ the *tree attached to* G. The model J is constructed as follows. Let J_0 be the instance C. Furthermore, for each guarded set $G = \{a_1, \dots, a_k\}$ in D and each copy $\Gamma = \{\alpha_1, \dots, \alpha_k\}$ of G in I^*, construct a new instance J_Γ that is isomorphic to the tree attached to G such that (i) the elements a_i of G are renamed to α_i in J_Γ, (ii) $\mathit{dom}(J_0) \cap \mathit{dom}(J_\Gamma) = \{\alpha_1, \dots, \alpha_k\}$, and (iii) $\Gamma \cap \Theta = \mathit{dom}(J_\Gamma) \cap \mathit{dom}(J_\Theta)$, for every copy Θ of G in I^*. The model J is the union of J_0 and all the J_Γ. If a guarded set X in J_Γ arises from renaming elements of a guarded set Y in $\mathit{chase}(D \restriction G, \Sigma)$, we also say that X is a copy of Y in J. Furthermore, the copies of D that are contained in I^* (i.e., in J_0) are also called copies in J. Observe that J is a model of I^* by construction. We show that it is a model of Σ. To this end, we show the following claim.

Claim 39. *Let $\bar{t}$ be a guarded tuple in J and let $q(\bar{x})$ be a guarded conjunctive query[9] over $\mathbf{S} \cup \mathit{sch}(\Sigma)$. Suppose $\bar{t}$ is a copy of $\bar{s}$ in J, where $\bar{s}$ is over $\mathit{dom}(\mathit{chase}(D, \Sigma))$ and $|\bar{t}| = |\bar{s}|$. Then $J \models q(\bar{t})$ iff $\mathit{chase}(D, \Sigma) \models q(\bar{s})$.*

Proof. Suppose $J \models q(\bar{t})$. Let $\{\bar{t}\} = \{\alpha_1, \dots, \alpha_k\}$ be a copy of $\{\bar{s}\} = \{a_1, \dots, a_k\}$ in J. Since $q(\bar{x})$ is guarded, there is a $\Gamma \supseteq (\{\alpha_1, \dots, \alpha_k\} \cap \mathit{dom}(J_0))$ such that $J_\Gamma \models q(\bar{t})$. Let $G \supseteq \{\bar{s}\}$ be the guarded set in D of which Γ is a copy in J_0. It clearly holds that $\mathit{chase}(D \restriction G, \Sigma) \models q(\bar{s})$, whence $\mathit{chase}(D, \Sigma) \models q(\bar{s})$ follows.

Suppose that $\mathit{chase}(D, \Sigma) \models q(\bar{s})$. Let $\bar{t} = \alpha_1, \dots, \alpha_k$ and $\bar{s} = a_1, \dots, a_k$ and suppose that $\alpha_i = [\pi_i]_{a_i}$ $(i = 1, \dots, k)$. The set $\{a_1, \dots, a_k\}$ is guarded in $\mathit{chase}(D, \Sigma)$. Hence, there is a guarded $G \supseteq \{a_1, \dots, a_k\} \cap \mathit{dom}(D)$ in D such that $\mathit{chase}(D \restriction G, \Sigma) \models q(\bar{s})$. We show that there is a $\Gamma \supseteq \{\alpha_1, \dots, \alpha_k\} \cap \mathit{dom}(I^*)$ which is a copy of G in I^*. Suppose $G = \{b_1, \dots, b_l\}$. Let $\pi = X_0 X_1 \cdots X_m$ be such that $[\pi]_{a_i} = [\pi_i]_{a_i}$ for all $i = 1, \dots, k$. For $i = 1, \dots, l$, define

$$\rho_i \coloneqq \pi \cdot (X_m \cup \{b_1\}) \cdot (X_m \cup \{b_1, b_2\}) \cdots (X_m \cup \{b_1, \dots, b_i\}).$$

Then b_i is represented at ρ_i. For $i = 1, \dots, l$, let $\beta_i \coloneqq [\rho_i]_{b_i}$. We claim that $\Gamma \coloneqq \{\beta_1, \dots, \beta_l\}$ is a copy of G in I^*. Let $R/s \in \mathbf{S}$ and suppose $I^* \models R([\rho_{i_1}]_{b_{i_1}}, \dots, [\rho_{i_s}]_{b_{i_s}})$. Then we immediately obtain $D \models R(b_{i_1}, \dots, b_{i_s})$. Conversely, if $D \models R(b_{i_1}, \dots, b_{i_s})$, let $\rho \coloneqq \rho_\ell$, where $\ell \coloneqq \max\{i_1, \dots, i_s\}$. Take any $j \in \{i_1, \dots, i_s\}$. It is easy to see that ρ and ρ_j are b_j-equivalent. Hence, $[\rho_j]_{b_j} = [\rho]_{b_j}$ and it follows that $I^* \models R([\rho_{i_1}]_{b_{i_1}}, \dots, [\rho_{i_s}]_{b_{i_s}})$, as required. It follows that Γ is a copy of G in I^* and so there is a structure J_Γ contained in J that is isomorphic to $\mathit{chase}(D \restriction G, \Sigma)$ with $b_1, \dots, b_l$ respectively renamed to $\beta_1, \dots, \beta_l$. Hence, $J \models q(\bar{t})$ as required. □

Now let $\sigma \colon \varphi(\bar{x}, \bar{y}) \to \exists \bar{z}\, \alpha(\bar{x}, \bar{z})$ be a guarded rule from Σ. Suppose that $J \models \exists \bar{y}\, \varphi(\bar{t}, \bar{y})$. Since every guarded tuple in J is a copy of some guarded tuple in $\mathit{chase}(D, \Sigma)$, there is an $\bar{s}$, of which $\bar{t}$ is a copy, such that $\mathit{chase}(D, \Sigma) \models \exists \bar{y}\, \varphi(\bar{s}, \bar{y})$. Since $\mathit{chase}(D, \Sigma)$ is a model of Σ, we know that $\mathit{chase}(D, \Sigma) \models \exists \bar{z}\, \alpha(\bar{s}, \bar{z})$. It follows that $J \models \exists \bar{z}\, \alpha(\bar{s}, \bar{z})$ by the above claim, as required. It remains to show that J is universal:

Claim 40. *J is universal.*

Proof. It suffices to show that J can be homomorphically mapped to $\mathit{chase}(I^*, \Sigma)$ via a homomorphism η. We let η_0 be the homomorphism that maps every element of J_0 to itself. It remains to treat the structures J_Γ. Consider a copy $\Gamma = \{\alpha_1, \dots, \alpha_k\}$ in I^* of a set $G = \{b_1, \dots, b_k\}$ which is guarded in D. It suffices to show that J_Γ can be mapped to $\mathit{chase}(I^*, \Sigma)$. To this end, it we show how to map $\mathit{chase}(D \restriction G, \Sigma)$ to $\mathit{chase}(I^*, \Sigma)$. We do so by induction on the number of rule applications of $\mathit{chase}(D \restriction G, \Sigma)$. For the base case, we map $D \restriction G$ to I^* as follows. Let $\eta^0_G(b_i) \coloneqq \alpha_i$, for $i = 1, \dots, k$. Suppose $D \restriction G \models R(b_{i_1}, \dots, b_{i_l})$ for some $R \in \mathbf{S}$ and $i_j \in \{1, \dots, k\}$, where $j = 1, \dots, l$. Recall that Γ is guarded in I^*. Reviewing the construction of I^*, it is easy to see that this holds iff $I^* \models R(\alpha_{i_1}, \dots, \alpha_{i_l})$. Hence, η^0_G is indeed a homomorphism from $D \restriction G$ to I^*. The induction step is obvious—we can easily obtain a homomorphism η^i_G that maps $\mathit{chase}^k(D \restriction G, \Sigma)$ to $\mathit{chase}(I^*, \Sigma)$. The desired homomorphism η_G is the union of the η^i_G $(i \geq 0)$. We then obtain a homomorphism η_Γ from η_G by appropriately renaming the elements from the domain of the latter as we did in the construction of J_Γ—which is nothing else than an isomorphic copy of $\mathit{chase}(D \restriction G, \Sigma)$. Furthermore, each of these homomorphisms maps each element of Γ to itself. The desired homomorphism η that witnesses that J is universal is the union of η_0 and the η_Γ. □

In order to prove $I^*, \Sigma \models q$, it remains to show that there is a homomorphism $\hat{\mu}$ that maps q to J. There are guarded sets $G_1, \dots, G_l$ in D such that μ can be understood to map q to $\mathit{chase}(\bigcup_{1 \leq i \leq l}(D \restriction G_i), \Sigma)$. By construction, we know that $G_1, \dots, G_l$ can be chosen in such a way that $G_\mu \subseteq \bigcup_{i=1}^{l} G_i$. Since Σ is guarded, μ can be understood to map q to $\bigcup_{1 \leq i \leq l} \mathit{chase}(D \restriction G_i, \Sigma)$—assuming that the labeled nulls occurring in these instances are mutually new. Let $\mathcal{C}_\mu \coloneqq \{\{\bar{b}_1\}, \dots, \{\bar{b}_k\}\}$. For every $X \in \mathcal{C}_\mu$, let $\Gamma_X \coloneqq \{\lambda_b \mid b \in X\}$. Notice that Γ_X is a copy of X in I^*. By construction, all the facts from q that are mapped via μ to $\mathit{chase}(D, \Sigma)$ and which have an element from $\mathit{dom}(D)$ in their image under μ are already mapped to $\bigcup_{X \in \mathcal{C}_\mu} \mathit{chase}(D \restriction X, \Sigma)$. For the other facts, the names of the constants in the databases do not matter.[10] Let $\Theta_1, \dots, \Theta_s$ be arbitrary copies of the sets $\{G_1, \dots, G_l\} \setminus \mathcal{C}_\mu$ in I^*. It follows that we can find our desired match $\hat{\mu}$ in the

[9] By a *guarded conjunctive query* we mean here a CQ that contains an atom that contains all the variables occurring in the CQ as argument.

[10] Here, it is of course essential to assume constant-free rules.

union of $\bigcup_{X\in\mathcal{C}_\mu} J_{\Gamma_X}$ and $\bigcup_{1\leq i\leq s} J_{\Theta_i}$. Notice that $\bigcup_{X\in\mathcal{C}_\mu} J_{\Gamma_X}$ is isomorphic to $\bigcup_{X\in\mathcal{C}_\mu} \mathit{chase}(D \restriction X, \Sigma)$ with each $b \in G_\mu$ represented by λ_b. □

Now the database I^* has the desired form with C being its core. However, I^* is infinite. Since $I^*, \Sigma \models q$ due to Lemma 38, by compactness, there is a finite $\hat{B} \subseteq I^*$ such that $\hat{B}, \Sigma \models q$. Consider a tree decomposition $\delta = (\mathcal{T}, (X_t)_{t\in T})$ witnessing that I^* is a C-tree. There is a maximum ℓ such that $\hat{B}$ contains all the subinstances induced by the bags of depth less or equal ℓ. Let $\hat{I}$ be the instance that actually contains *all* the subinstances induced by the bags of level up to ℓ. Hence, $\hat{I}$ is itself a C-tree and $\hat{I}, \Sigma \models q$, since $\hat{B} \subseteq \hat{I}$.

Now there is a natural homomorphism mapping $\hat{I}$ to D: we simply specify $[\pi]_a \mapsto a$ for all $a \in \mathit{dom}(D)$. The instance $\hat{I}$ is the one we are looking for. □

Proof of Lemma 22

One can naturally encode instances of bounded tree-width into trees over a finite alphabet such that the alphabet's size depends only on the tree-width. Our goal here is to appropriately encode C-trees in order to make them accessible to tree automata techniques. Since the tree-width of a C-tree over $\mathbf{S}$ depends only on the size of $\mathit{dom}(C)$ and the maximum arity of $\mathbf{S}$, the alphabet of the encoding will depend on the same.

Labeled trees. Let Γ be an alphabet and $(\mathbb{N}\setminus\{0\})^*$ be the set of finite sequences of positive integers, including the empty sequence ε.[11] Let us recall that a Γ-*labeled tree* is a pair $t = (T, \mu)$, where $\mu\colon T \to \Gamma$ is the *labeling function* and $T \subseteq (\mathbb{N}\setminus\{0\})^*$ is closed under prefixes, i.e., $x \cdot i \in T$ implies $x \in T$, for all $x \in (\mathbb{N}\cup\{0\})^*$ and $i \in (\mathbb{N}\cup\{0\})$. The elements contained in T identify the *nodes* of t. For $i \in \mathbb{N}\setminus\{0\}$, nodes of the form $x \cdot i \in T$ are the *children* of x. A *path* of length n in T from x to y is a sequence of nodes $x = x_1, \ldots, x_n = y$ such that x_{i+1} is a child of x_i. A *branch* is a path from the root to a leaf node. For $x \in T$, we set $x \cdot i \cdot -1 \coloneqq x$, for all $i \in \mathbb{N}$, and $x \cdot 0 \coloneqq x$—notice that $\varepsilon \cdot -1$ is not defined.

Encoding. Let $l \geq 0$ and fix a schema $\mathbf{S}$. Let $U_{\mathbf{S},l}$ be the disjoint union of two sets C_l and $T_{\mathbf{S}}$, respectively containing l and $2 \cdot \mathit{ar}(\mathbf{S})$ elements. The elements from $U_{\mathbf{S},l}$ will be called *names*. Elements from the set C_l will describe core elements, while those of $T_{\mathbf{S}}$ will describe the others. Furthermore, neighboring nodes may describe overlapping pieces of the instance. In particular, if one name is used in neighboring nodes, this means that the name at hand refers to the same element—this is why we use $2w$ elements for the non-root bags. Let $\mathbb{K}_{\mathbf{S},l}$ be the finite schema capturing the following information:

- For all $a \in U_{\mathbf{S},l}$, there is a unary relation $D_a \in \mathbb{K}_{\mathbf{S},l}$.
- For all $a \in C_l$, there is a unary relation $C_a \in \mathbb{K}_{\mathbf{S},l}$.
- For each $R \in \mathbf{S}$ and every n-tuple $\bar{a} \in U^n_{\mathbf{S},l}$, there is a unary relation $R_{\bar{a}} \in \mathbb{K}_{\mathbf{S},l}$.

Let $\Gamma_{\mathbf{S},l} \coloneqq 2^{\mathbb{K}_{\mathbf{S},l}}$ be an alphabet and suppose that D is a (finite) C-tree over $\mathbf{S}$ such that $|\mathit{dom}(C)| \leq l$. Consider a tree decomposition $\delta = (\mathcal{T}, (X_t)_{t\in T})$ witnessing that D is indeed a C-tree and let ε be the root of $\mathcal{T}$. Fix a function $f\colon \mathit{dom}(D) \to U_{\mathbf{S},l}$ such that (i) $f \restriction \mathit{dom}(C)$ is injective and (ii) different elements that occur in neighboring bags of δ are always assigned different names from $U_{\mathbf{S},l}$. Using f, we can encode D and δ into a $\Gamma_{\mathbf{S},l}$-labeled tree $t = (\hat{T}, \mu)$ such that each node from $\mathcal{T}$ corresponds to exactly one node in $\hat{T}$ and vice versa. For a node v from $\mathcal{T}$, we denote the corresponding node of T by $\hat{v}$ in the following and vice versa. In this light, the symbols from $\mathbb{K}_{\mathbf{S},l}$ have the following intended meaning:

- $D_a \in \mu(\hat{v})$ means that a is used as a name for some element of the bag X_v.
- $C_a \in \mu(\hat{v})$ indicates that a is used as name for an element of the bag X_v that also occurs in X_ε, i.e., a names an element from the core of D.
- $R_{\bar{a}} \in \mu(\hat{v})$ indicates that R holds in D for the elements named by $\bar{a}$ in bag X_v.

Under certain assumptions, we can decode a $\Gamma_{\mathbf{S},l}$-labeled tree $t = (T, \mu)$ into a C-tree whose width is bounded by $\mathit{ar}(\mathbf{S}) - 1$. Let $\mathrm{names}(v) \coloneqq \{a \mid D_a \in \mu(v)\}$. We say that t is *consistent*, if it satisfies the following properties:

1. For all nodes v it holds that $|\mathrm{names}(v)| \leq \mathit{ar}(\mathbf{S})$, except for the root whose number of names are accordingly bounded by l. Furthermore, $\mathrm{names}(\varepsilon) \subseteq C_l$.
2. For all $R_{\bar{a}} \in \mathbb{K}_{\mathbf{S},l}$ and all $v \in T$ it holds that $R_{\bar{a}} \in \mu(v)$ implies that $\{\bar{a}\} \subseteq \mathrm{names}(v)$.
3. For all $a \in C_l$ and all $v \in T$ it holds that $D_a \in \mu(v)$ iff $C_a \in \mu(v)$.
4. If $C_a \in \mu(v)$, then $C_a \in \mu(w)$ for all $w \in T$ on the unique shortest path between v and the root.
5. For all nodes $v \neq \varepsilon$, there is an $R_{\bar{a}} \in \mathbb{K}_{\mathbf{S},l}$ and a node w such that $R_{\bar{a}} \in \mu(w)$, $\mathrm{names}(v) \subseteq \{\bar{a}\}$, and, for all $b \in \mathrm{names}(v)$, v and w are b-connected.

[11] We specify that 0 is included in $\mathbb{N}$ as well.

Decoding trees. Suppose now that t is consistent. We show how we can decode t into a database $[\![t]\!]$ which is a C-tree whose diameter is bounded by l. Let a be a name used in t. We say that two nodes v, w of t are *a-equivalent* if $D_a \in \mu(u)$ for all nodes u on the unique shortest path between v and w. Clearly, a-equivalence defines an equivalence relation and we let $[v]_a \coloneqq \{(w, a) \mid w \text{ is } a\text{-equivalent to } v\}$ and $[v]^*_a \coloneqq \{w \mid (w, a) \in [v]_a\}$. The domain of $[\![t]\!]$ is the set $\{[v]_a \mid v \in T, a \in \mu(v)\}$ and, for $R/n \in \mathbf{S}$, we define

$$[\![t]\!] \models R([v_1]_{a_1}, \dots, [v_n]_{a_n}) \iff_{df} \text{ there is some } v \in [v_1]^*_{a_1} \cap \dots \cap [v_n]^*_{a_n} \text{ such that } R_{a_1,\dots,a_n} \in \mu(v).$$

LEMMA 41. *Let t be a consistent $\Gamma_{\mathbf{S},l}$-labeled tree with root node ε. Then $[\![t]\!]$ is well-defined and a C-tree over $\mathbf{S}$, where C is the subinstance of $[\![t]\!]$ induced by the set $\{[\varepsilon]_a \mid a \in \text{names}(\varepsilon)\}$. Moreover, $|dom(C)|$ is bounded by l.*

PROOF. Let $t = (T, \mu)$ be a consistent, $\Gamma_{\mathbf{S},l}$-labeled tree. The fact that $[\![t]\!]$ is well-defined is left to the reader. We are going to construct an appropriate decomposition $\delta = (\mathcal{T}, (X_t)_{t \in T})$ for $[\![t]\!]$. The tree $\mathcal{T}$ has the same structure as t. Furthermore, for $v \in T$, we set $X_v \coloneqq \{[v]_a \mid a \in \text{names}(v)\}$. We need to show that δ is indeed a tree decomposition that satisfies the desired properties.

Let $[v]_a \in dom([\![t]\!])$ and consider two nodes $v_1, v_2 \in T$ such that $[v]_a \in X_{v_1}$ and $[v]_a \in X_{v_2}$. Then $v_1, v_2 \in [v]_a$ and so v_1 and v_2 are a-connected. Hence, $w \in [v]_a$ for all $w \in T$ which lie on the unique shortest path between v_1 and v_2. Since $a \in \text{names}(w)$ for all such w, it follows that $[v]_a \in X_w$, and so $[v]_a$ is contained in all bags on the unique path between v_1 and v_2. Suppose $[\![t]\!] \models R([v_1]_{a_1}, \dots, [v_n]_{a_n})$. Then there is a $v \in [v_1]^*_{a_1} \cap \dots \cap [v_n]^*_{a_n}$ such that $R_{a_1,\dots,a_n} \in \mu(v)$. By consistency, $\{a_1, \dots, a_n\} \subseteq \text{names}(v)$. Moreover, we know that $[v_i]_{a_i} = [v]_{a_i}$, for $i = 1, \dots, n$. It follows that $\{[v_1]_{a_1}, \dots, [v_n]_{a_n}\} \subseteq X_v$. Now let $v \in T \setminus \{\varepsilon\}$. By consistency, there is an $R_{a_1,\dots,a_n} \in \mathbb{K}_{\mathbf{S},l}$ and a $w \in T$ such that $\text{names}(v) \coloneqq \{a_{i_1}, \dots, a_{i_s}\} \subseteq \{a_1, \dots, a_n\} \subseteq \text{names}(w)$, $R_{a_1,\dots,a_n} \in \mu(w)$, and v and w are b_{i_j}-connected for $j = 1, \dots, s$. By construction, $X_v = \{[v]_{a_{i_1}}, \dots, [v]_{a_{i_s}}\}$ and $\{[w]_{a_1}, \dots, [w]_{a_n}\} \subseteq X_w$. The claim follows now since $[v]_{a_{i_j}} = [w]_{a_{i_j}}$ for $j = 1, \dots, s$. It is immediate that $|dom(C)|$ is bounded by l. □

Notation. Given a consistent $\Gamma_{\mathbf{S},l}$-labeled tree $t = (T, \mu)$ and a label $\rho \in \mu(T)$, in order to ease notation we often regard ρ as a database consisting of the facts $\{R(\bar{a}) \mid R_{\bar{a}} \in \rho\}$. Furthermore, we let $\text{names}(\rho) \coloneqq \{a \mid D_a \in \rho\}$.

PROOF OF LEMMA 22. The lemma is an easy consequence of Lemma 41 and the fact that, when encoding a C-tree D over $\mathbf{S}$, together with a tree decomposition witnessing that D is a C-tree, into a consistent $\Gamma_{\mathbf{S},l}$-labeled tree t, then $[\![t]\!]$ and D are isomorphic. □

Roughly, Lemma 22 states that containment among OMQs from $(\mathbb{G}, \mathbb{BCQ})$ can be semantically characterized via the decodings of consistent $\Gamma_{\mathbf{S},l}$-labeled trees. This makes the problem of deciding containment amenable to tree automata techniques.

Proof of Lemma 23

Before proceeding to the proof of Lemma 23, we first introduce the relevant automata model.

Automata Techniques

For a set of propositional variables X, we denote by $\mathbb{B}^+(X)$ the set of Boolean formulas using variables from X, the connectives $\wedge, \vee$, and the constants $\mathsf{true}, \mathsf{false}$. Let us now introduce our automata model.

Definition 10. A *two-way alternating parity automaton (2WAPA) on trees* is a tuple $\mathfrak{A} = (S, \Gamma, \delta, s_0, \Omega)$, where S is a finite set of *states*, Γ an alphabet (the *input alphabet* of $\mathfrak{A}$), $\delta \colon S \times \Gamma \to \mathbb{B}^+(\mathsf{tran}(\mathfrak{A}))$ the *transition function*, where we set $\mathsf{tran}(\mathfrak{A}) \coloneqq \{\langle\alpha\rangle s, [\alpha]s \mid s \in S, \alpha \in \{-1, 0, *\}\}$, $s_0 \in S$ the *initial state*, and $\Omega \colon S \to \mathbb{N}$ the *parity condition* that assigns to each $s \in S$ a *priority* $\Omega(s)$. Elements from $\mathsf{tran}(\mathfrak{A})$ are called *transitions*. ■

Intuitively, a transition of the form $\langle 0\rangle s$ means that a copy of the automaton should change to state s and stay at the current node. A transition of the form $\langle -1\rangle s$ means that a copy should be sent to the parent node, which is then required to exist, and proceed in state s, while one of the form $\langle *\rangle s$ means that a copy of the automaton that assumes state s is sent to some child node. The transition $[0]s$ means the same as $\langle 0\rangle s$, while $[-1]s$ means that a copy of the automaton that assumes state s should be sent to the parent node which is there not required to exist at all. Likewise, $[*]s$ means that a copy of the automaton assuming state s should be sent to all child nodes.

Notation. We write $\Diamond s$ for $\bigvee\{\langle\alpha\rangle s \mid \langle\alpha\rangle s \in \mathsf{tran}(\mathfrak{A}), s \in S\}$, $\Box s$ for $\bigwedge\{[\alpha]s \mid [\alpha]s \in \mathsf{tran}(\mathfrak{A}), s \in S\}$, and simply s for $\langle 0\rangle s$. Furthermore, for $\alpha \in \mathbb{N} \cup \{-1, *\}$, we define

$$T_\alpha(x) \coloneqq \begin{cases} \{x \cdot \alpha\}, & \text{if } \alpha = -1 \text{ and } x \cdot \alpha \in T, \\ \{x \cdot i \mid x \cdot i \in T, i \in \mathbb{N} \setminus \{0\}\}, & \text{if } \alpha = *, \\ \varnothing, & \text{otherwise.} \end{cases}$$

Definition 11. A *run* of a 2WAPA $\mathfrak{A} = (S, \Gamma, \delta, s_0, \Omega)$ on a Γ-labeled tree (T, η) is a $T \times S$-labeled tree (T_r, η_r) such that the following holds:

1. $\eta_r(\varepsilon) = (\varepsilon, s_0)$,
2. if $y \in T_r$, $\eta_r(y) = (x, s)$, and $\delta(s, \eta(x)) = \varphi$, then there is an $I \subseteq \mathsf{tran}(\mathbb{A})$ such that $I \models \varphi$ holds and the following conditions are satisfied:
   - If $\langle\alpha\rangle s' \in I$ then there is a node $x' \in T_\alpha(x)$ and a child node $y' \in T_r$ of y such that $\eta_r(y') = (x', s')$.
   - If $[\alpha]s' \in I$ then for all $x' \in T_\alpha(x)$, there is a child node $y' \in T_r$ of y such that $\eta_r(y') = (x', s')$.

We say that a run (T_r, η_r) is *accepting* on $\mathfrak{A}$, if on all infinite paths $(\varepsilon, s_0), (x_1, s_1), (x_2, s_2), \dots$ in T_r, the maximum priority among $\Omega(s_0), \Omega(s_1), \Omega(s_2), \dots$ that appears infinitely often is even. $\mathfrak{A}$ *accepts* a Γ-labeled tree (T, η), if there is an accepting run on (T, η). We denote by $\mathcal{L}(\mathfrak{A})$ the set of Γ-labeled trees $\mathfrak{A}$ accepts, i.e., the *language accepted by* $\mathfrak{A}$. ∎

Remark. The automaton model defined above resembles that in [58]. However, we explicitly provide transitions that allow the automaton move to the parent node, while the model defined in [58] provides transitions for moving to *some* neighboring node, including the parent node. Therefore, the automata in [58] offer transitions of the form s, $\Diamond s$, and $\Box s$ with their intended meaning as defined above. Using techniques as employed in [57, 58], for a 2WAPA $\mathfrak{A}$, one can show that the problem of deciding whether $\mathcal{L}(\mathfrak{A}) = \varnothing$ is feasible in exponential time with respect to the number of states of $\mathfrak{A}$ and in polynomial time with respect to the size of the input alphabet of $\mathfrak{A}$.

Proof of Lemma 23. We only give an intuitive explanation for the construction of the desired 2WAPA. To check whether a $\Gamma_{\mathbf{S},l}$-labeled tree is consistent, we can check each condition for consistency separately by a dedicated 2WAPA and then take the intersection of all of them. Most of the consistency conditions are easy to check. We give here a more detailed verbal explanation for condition (5). A 2WAPA checking this condition can be constructed as follows. At the beginning of its run, the automaton branches universally to all nodes (except the root) in a state whose intended purpose is to find appropriate guards in the input tree for the names available at the current node. To this end, the automaton has to do a reachability analysis on the input tree and store, using exponentially many states in $ar(\mathbf{S})$, the tuple it seeks to guard. By a *guard* for the node v here, we mean a node w with an $R_{\bar{a}} \in \mu(w)$ such that (i) $\{\bar{a}\}$ contains all the names present at w and (ii) is b-connected to v for all $b \in \text{names}(v)$. Notice that such a reachability analysis can be easily performed once we have the means to store the information contained in names(v) in a single state. This is, however, possible since for this task we need somewhat $O((ar(\mathbf{S})+l)^{ar(\mathbf{S})})$ states, i.e., polynomially many in the size of $\Gamma_{\mathbf{S},l}$. □

Proof of Lemma 24

We first need to introduce some additional auxiliary notions.

Strictly Acyclic Queries

Let q be a CQ over a schema $\mathbf{S}$. We denote by free(q) the free variables of q; the same notation is used for first-order formulas in general. We can naturally view q as an instance $[q]$ whose domain is the set of variables of q and contains the body atoms of q as facts. In the following, we will often overload notation and write q for both the query q and the instance $[q]$. The notions of tree-width, acyclicity, etc. then immediately extend to CQs. Given a tree decomposition δ of q (i.e., of $[q]$), we say that δ is *strict*, if some bag of δ contains all variables that are free in q (cf. also [38]). Accordingly, q is called *strictly acyclic* if it has a guarded tree decomposition that is strict.

Strictly acyclic queries have the convenient property to be equivalent to *guarded formulas* of a special form. Recall that the set of *guarded formulas* over a schema $\mathbf{S}$ is built inductively by including all atomic formulas, relativizing quantifiers by atomic formulas, and closing under Boolean connectives. More precisely, all quantifier occurrences have one of the forms

$$\forall \bar{y}\,(\alpha(\bar{x}, \bar{y}) \to \varphi) \quad \text{and} \quad \exists \bar{y}\,(\alpha(\bar{x}, \bar{y}) \wedge \varphi),$$

such that the free variables of φ are among $\{\bar{x}, \bar{y}\}$.

We are interested in the guarded formulas that are build up using conjunction and existential quantification; we restrict ourselves to such formulas in the following. We call a formula from this class *strictly guarded*, if it is of the form $\exists \bar{y}\,(\alpha(\bar{x}, \bar{y}) \wedge \varphi)$. We explicitly include the case where $\bar{y}$ is the empty sequence of variables, i.e., if to formulas of the form $\alpha(\bar{x}) \wedge \varphi$ with free$(\varphi) \subseteq \{\bar{x}\}$. Notice that every guarded sentence φ (i.e., a formula having no free variables) is strictly guarded, since it is equivalent to $\exists y\,(y = y \wedge \varphi)$. Furthermore, notice that every usual guarded formula that uses only existential quantifiers and conjunction is equivalent to a conjunction of strictly guarded formulas. The following lemma is proved in [38].

Lemma 42. *Every strictly acyclic CQ can be rewritten in polynomial time into an equivalent strictly guarded formula that is built up using conjunction and existential quantification only. The converse holds as well.*

Squid Decompositions

Let q be a BCQ over a schema $\mathbf{S}$ having n body atoms. An $\mathbf{S}$-*cover* of q is a BCQ q^+ that contains all the atoms from q and may additionally contain $2n$ other body atoms over $\mathbf{S}$. It is pretty straightforward that, for an $\mathbf{S}$-instance I, it holds that $I \models q$ iff there is an $\mathbf{S}$-cover q^+ of q such that $I \models q^+$.

Definition 12. Let I be an instance. For $V \subseteq dom(I)$, we say that I is $[V]$-*acyclic*, if it has a guarded tree decomposition that omits V. ■

Definition 13. Let q be a BCQ over $\mathbf{S}$. A *squid decomposition* of q is a tuple $\delta = (q^+, \mu, H, T, V)$, where q^+ is an $\mathbf{S}$-cover of q, $\mu\colon var(q^+) \to var(q^+)$ a mapping, $V \subseteq var(\mu(q^+))$, and (H, T) a partition of the atoms $\mu(q^+)$ such that

- H is the set of atoms of $\mu(q^+)$ induced by V,
- $T = \mu(q^+) \setminus H$ and T is $[V]$-acyclic. ■

Intuitively, a squid decomposition specifies a way how a BCQ can be mapped to an instance that contains some "cyclic parts"—the set H specifies those atoms that are mapped to such cyclic parts, while A declares those atoms that are mapped to the acyclic parts of the instance at hand. We will make this more precise in Lemma 43 below, where we analyze matches in C-trees.

Given a CQ q and a set of variables $V \subseteq var(q)$, the V-*reduct* of q, denoted q^V, is the conjunctive query that arises from q by dropping all the existential quantifiers that bind variables in V.

LEMMA 43. *Let J be a C-tree over $\mathbf{S}$ and q a BCQ over $\mathbf{S}$. Let $(\mathcal{T}, (X_t)_{t\in T})$ be a witnessing tree decomposition of J. It holds that $J \models q$ iff there is a squid decomposition $\delta = (q^+, \mu, H, A, V \coloneqq \{\bar{x}\})$ of q and a homomorphism $\eta\colon \mu(q^+) \to J$ such that*

1. *$C \models H$ is witnessed by η,*
2. *$\bigcup_{\varepsilon \prec v} J(v) \models A^V(\eta(\bar{x}))$ is witnessed by η, and*
3. *there are strictly guarded formulas $\varphi_1, \ldots, \varphi_l$ such that $A^V(\bar{x}) \equiv \varphi_1 \wedge \cdots \wedge \varphi_l$.*

PROOF. For the direction from right to left, consider such a given squid decomposition δ and a homomorphism η as in the hypothesis of the lemma. It is immediate that $\eta \circ \mu$ is a homomorphism mapping q^+ to J. Since q^+ is an $\mathbf{S}$-cover of q, we obtain $J \models q$ as required.

For the other direction, suppose that $J \models q$ is witnessed by a homomorphism θ. For each $v \in T \setminus \{\varepsilon\}$, let β_v be an atom of J such that $J(v) \models \beta_v$ and β_v contains all domain elements from $J(v)$ as arguments. Notice that the β_v exist, since δ is guarded except for $\{\varepsilon\}$. Since θ maps q to J, for each atom α of q, there is a node v_α such that $\theta(\alpha) \in J(v_\alpha)$. Let W be the set of all these nodes and their closure under greatest lower bounds with respect to $\preceq$, excluding the root node ε of $\mathcal{T}$. Consider the set of atoms $Q^+ \coloneqq \theta(q) \cup \{\beta_v \mid v \in W\}$. Notice that at least half of the nodes of W are of the form v_α—hence, $|Q^+| \leq 3|q|$. Let q^+ be a BCQ constructed as follows. Take the conjunction of q and for each $\beta_v(a_1, \ldots, a_n)$ $(v \in W)$, add an atom $\beta_v(x_1, \ldots, x_n)$, where each x_i is a newly chosen variable. Then q^+ is obviously an $\mathbf{S}$-cover of q. Furthermore, by construction, there is a mapping $\mu\colon var(q^+) \to var(q^+)$ and an isomorphism $\eta\colon \mu(q^+) \to Q^+$ such that $(\eta \circ \mu)(q^+) = Q^+$. Now let H be the greatest set of atoms of $\mu(q^+)$ such that $\eta(H) \subseteq J(\varepsilon)$. Moreover, let $V \coloneqq var(H)$ and $A \coloneqq \mu(q^+) \setminus H$. We claim that $\delta \coloneqq (q^+, \mu, H, A, V)$ is a squid decomposition of q that satisfies together with η the points mentioned in the statement of the lemma.

To see that δ is a squid decomposition of q, the only nontrivial point to prove is that A is indeed $[V]$-acyclic. We will prove this below in the course of establishing the third item.

The first two items are immediate by construction. We prove the third item. Suppose $V = \{\bar{x}\}$ and consider the V-reduct $A^V(\bar{x})$ of A. By construction, the atoms $\eta(A)$ are contained in $\bigcup_{\varepsilon \prec v} J(v)$. Now the set W together with the order $\preceq_\mathcal{T}$ gives rise to a forest consisting of trees $\mathcal{T}_1, \ldots, \mathcal{T}_l$ whose roots are descendants of ε, i.e., the root of $\mathcal{T}$ (recall that ε is not contained in W). Moreover, we have (i) $\bigcup_{i=1}^{l} T_i = W$, (ii) $T_i \cap T_j = \varnothing$, for $i \neq j$, and (iii) $\bigcup_{v\in T_i} X_v \cap \bigcup_{v\in T_j} X_v \subseteq dom(C)$, for $i \neq j$.

For $v \in T$, let $Q^+(v) \coloneqq \{\alpha \in Q^+ \mid J(v) \models \alpha\}$ and, for $i = 1, \ldots, l$, let $Q^+(\mathcal{T}_i)$ be the set of atoms $\bigcup_{v\in T_i} Q^+(v)$. Now it is easy to check using the facts stated before that each $Q^+(\mathcal{T}_i)$ is acyclic and, hence, so is $\eta^{-1}(Q^+(\mathcal{T}_i))$. Furthermore, denoting by ε_i the root of $\mathcal{T}_i$, it holds that $dom(Q^+(\mathcal{T}_i)) \cap \eta(V) \subseteq dom(Q^+(\varepsilon_i))$—indeed, if $a \in dom(Q^+(v)) \cap \eta(V)$ for some $v \succeq \varepsilon_i$, then, since $\varepsilon_i \succ \varepsilon$ and $a \in X_\varepsilon$, it must be the case that $a \in dom(Q^+(\varepsilon_i))$ by connectivity. It follows that the V-reduct of $\eta^{-1}(Q^+(\mathcal{T}_i))$ (viewed as Boolean query), henceforth denoted $q^+_{\mathcal{T}_i}$, is strictly acyclic and is therefore equivalent to a strictly guarded formula φ_i. Hence, the query $A^V(\bar{x})$ is equivalent to $\bigwedge_{i=1}^{l} \varphi_i$. Moreover, it follows that A itself is $[V]$-acyclic—notice that $A \equiv \exists\bar{x} \bigwedge_{i=1}^{l} q^+_{\mathcal{T}_i}$ and that $dom(Q^+(\mathcal{T}_i)) \cap dom(Q^+(\mathcal{T}_j)) \subseteq \eta(V)$, for $i \neq j$. Hence, $var(q^+_{\mathcal{T}_i}) \cap var(q^+_{\mathcal{T}_j}) \subseteq V$, for $i \neq j$. The claim now follows since every $q^+_{\mathcal{T}_i}$ is acyclic. □

Derivation trees

Let D be an $\mathbf{S}$-database and Σ a set of guarded rules. Let $q_0(\bar{x})$ be a strictly acyclic query whose free variables are exactly those from $\bar{x} \coloneqq x_1, \ldots, x_n$ and let $\bar{a} \coloneqq a_1, \ldots, a_n$ be a tuple from $dom(D)$. A *derivation tree for $(\bar{a}, q_0(\bar{x}))$ with respect to D and Σ* is a finite tree $\mathcal{T}$ whose nodes are labeled via a function μ with pairs of the form $(b_1, \ldots, b_k; q(y_1, \ldots, y_k))$, where $b_1, \ldots, b_k$ are constants from $dom(D)$ and $q(y_1, \ldots, y_k)$ is a strictly acyclic query over $\mathbf{S} \cup sch(\Sigma)$ having exactly $y_1, \ldots, y_k$ free, such that the following conditions are satisfied:

1. $\mu(\varepsilon) = (\bar{a}, q_0(\bar{x}))$, where ε is the root node of $\mathcal{T}$.

2. If $\mu(v) = (c_1, \ldots, c_m; q(z_1, \ldots, z_m))$ for some node v, then one of the following conditions holds (let $\bar{c} \coloneqq c_1, \ldots, c_m$ and $\bar{z} \coloneqq z_1, \ldots, z_m$):

   (a) v is a leaf node and $q(\bar{z}) \equiv \beta(\bar{z})$, for some atomic formula $\beta(\bar{z})$ such that $D \models \beta(\bar{c})$.

   (b) The node v has a successor labeled by $(\bar{c}, \bar{b}; p(\bar{z}, \bar{y}))$ and it holds that
   $$\Sigma \models \forall \bar{z}, \bar{y}\, (p(\bar{z}, \bar{y}) \rightarrow q(\bar{z})).$$

   (c) The query $q(\bar{z})$ is logically equivalent to $q_1(z_{i_{1,1}}, \ldots, z_{i_{1,k_1}}) \wedge \cdots \wedge q_l(z_{i_{l,1}}, \ldots, z_{i_{l,k_l}})$ and v has l successors $v_1, \ldots, v_l$ respectively labeled by $(c_{i_{1,1}}, \ldots, c_{i_{1,k_1}}; q_1(z_{i_{1,1}}, \ldots, z_{i_{1,k_i}})), \ldots, (c_{i_{l,1}}, \ldots, c_{i_{l,k_l}}; q_l(z_{i_{l,1}}, \ldots, z_{i_{1,k_l}}))$.

LEMMA 44. *Let $\alpha(x_1, \ldots, x_n)$ be an atomic formula. Then $D, \Sigma \models \alpha(a_1, \ldots, a_n)$ iff there is a derivation tree for $(a_1, \ldots, a_n; \alpha(x_1, \ldots, x_n))$ with respect to D and Σ.*

PROOF (SKETCH). Let $\bar{a} \coloneqq a_1, \ldots, a_n$ and $\bar{x} \coloneqq x_1, \ldots, x_n$. The direction from right to left is an easy induction on the construction of the derivation tree. We sketch the other direction. Consider the guarded chase forest $(\mathcal{F}, \eta)$ for D and Σ, where η is a function labeling the nodes and edges of $\mathcal{F}$. We construct a derivation tree for $(\bar{a}, \alpha(\bar{x}))$ by induction on the number of chase steps required to derive $\alpha(\bar{a})$ from D and Σ.

For the base case, if $D \models \alpha(\bar{a})$, the claim is obvious since we can apply rule 2.(a). Assume that $\alpha(\bar{a})$ is derived using a rule

$$\sigma\colon \quad \beta_0(\bar{x}, \bar{y}), \beta_1, \ldots, \beta_k \rightarrow \alpha(\bar{x}),$$

and a homomorphism μ such that $\mu(\bar{x}) = \bar{a}$, where $\beta_0(\bar{x}, \bar{y})$ is the guard of σ. If $\mu(\{\bar{x}, \bar{y}\}) \subseteq \mathit{dom}(D)$, the result immediately follows by the induction hypothesis. Otherwise, the image of $\beta_0(\bar{x}, \bar{y})$ under μ contains some labeled nulls as arguments. Assume that all the $\beta_1, \ldots, \beta_k$ contain nulls as their arguments—for those that do not, the induction hypothesis would yield appropriate derivation trees again. Notice that all the nulls occurring in $\beta_1, \ldots, \beta_k$ appear in $\mu(\{\bar{y}\})$. By construction of $\mathcal{F}$, there is a node v_0 that is an ancestor of the nodes having the atoms $\mu(\beta_0), \mu(\beta_1), \ldots, \mu(\beta_k)$ as labels and which has a label of the form $\beta_0(\bar{a}, \bar{b})$ which contains no nulls at all as arguments. There is a corresponding atomic formula $\gamma_0(\bar{x}, \bar{z})$ whose image under an appropriate homomorphism equals $\beta_0(\bar{a}, \bar{b})$. Furthermore, there are atoms $\gamma_1, \ldots, \gamma_l$ such that $\mathit{dom}(\{\gamma_1, \ldots, \gamma_l\}) \subseteq \{\bar{a}, \bar{b}\}$ and

$$\Sigma \models \beta_0(\bar{a}, \bar{b}) \wedge \gamma_1 \wedge \cdots \wedge \gamma_l \rightarrow \exists \bar{y}\, (\beta_0(\bar{a}, \bar{y}) \wedge \beta_1 \wedge \cdots \wedge \beta_k).$$

Now regard the γ_i $(i = 1, \ldots, l)$ as atomic formulas with free variables among $\{\bar{x}, \bar{z}\}$. The formula $p(\bar{x}, \bar{z}) \coloneqq \gamma_0(\bar{x}, \bar{z}) \wedge \gamma_1 \wedge \cdots \wedge \gamma_l$ is then a strictly acyclic query that satisfies $\Sigma \models \forall \bar{x}, \bar{z}\, (p(\bar{x}, \bar{z}) \rightarrow \alpha(\bar{x}))$. An application of rule 2.(b) then requires us to find a derivation tree for $(\bar{a}, \bar{b}; p(\bar{x}, \bar{y}))$, whence an application of rule 2.(c) reduces this task to finding derivation trees for the atoms $\gamma_0, \gamma_1, \ldots, \gamma_l$ and their corresponding tuples of constants. These trees exist by induction hypothesis and we can simply concatenate them appropriately in order to arrive at a derivation tree for $(\bar{a}, \alpha(\bar{x}))$. □

Given a guarded formula $\varphi(\bar{x})$ built up from conjunctions and existential quantification, we define the *nesting depth* of $\varphi(\bar{x})$, denoted $\mathrm{nd}(\varphi(\bar{x}))$, inductively:

- If $\varphi(\bar{x})$ is an atomic formula, then $\mathrm{nd}(\varphi(\bar{x})) \coloneqq 0$.
- If $\varphi(\bar{x}) = (\psi_1 \wedge \psi_2)$, then $\mathrm{nd}(\varphi(\bar{x})) \coloneqq \max\{\mathrm{nd}(\psi_1), \mathrm{nd}(\psi_2)\}$.
- If $\varphi(\bar{x}) = \exists \bar{y}\, (\alpha(\bar{x}, \bar{y}) \wedge \psi)$ and $\bar{y} \neq \varnothing$, then $\mathrm{nd}(\varphi(\bar{x})) \coloneqq \mathrm{nd}(\psi) + 1$.

LEMMA 45. *Let D be a database, Σ a set of guarded rules, and $q(\bar{x})$ a strictly acyclic conjunctive query. Then $D, \Sigma \models q(\bar{a})$ iff there is a derivation tree for $(\bar{a}, q(\bar{x}))$ with respect to D and Σ.*

PROOF (SKETCH). We again sketch only the direction from left to right. Let $\varphi(\bar{x})$ be the strictly guarded formula corresponding to $q(\bar{x})$. We proceed by induction on the nesting depth of $\varphi(\bar{x})$. If $\mathrm{nd}(\varphi(\bar{x})) = 0$, then $\varphi(\bar{x})$ is quantifier free and thus a conjunction of atoms $\alpha_0(\bar{x}) \wedge \alpha_1 \wedge \cdots \wedge \alpha_k$, where $\mathit{var}(\alpha_i) \subseteq \{\bar{x}\}$ for $i = 1, \ldots, k$. An application of rule 2.(c) reduces the problem of building a derivation tree for $(\bar{a}, \varphi(\bar{x}))$ to the problem of building corresponding trees for the α_i and their corresponding constants from $\bar{a}$. The existence of these trees is guaranteed by Lemma 44.

Now suppose that $\mathrm{nd}(\varphi(\bar{x})) = n + 1$. Let $\varphi(\bar{x}) = \exists \bar{y}\, (\alpha(\bar{x}, \bar{y}) \wedge \psi)$ and $\bar{y} \coloneqq y_1, \ldots, y_k$. Assume, without loss of generality, that all the bound variables from $\varphi(\bar{x})$ are pairwise distinct. In the following, we will describe how to construct a derivation tree for $(\bar{a}, q(\bar{x}))$. If $D, \Sigma \models q(\bar{a})$, then there is a homomorphism μ mapping each atom of $q(\bar{x})$ to $\mathit{chase}(D, \Sigma)$ such that $\mu(\bar{x}) = \bar{a}$. Furthermore, μ maps each atom of $q(\bar{x})$ to a node of the guarded chase forest $\mathcal{F}$ of D and Σ. Let $\alpha_\mu(\bar{a}, \lambda_1, \ldots, \lambda_k)$ denote the atom labeling the node of $\mathcal{F}$ where $\alpha(\bar{x}, \bar{y})$ is mapped to via μ. Let $\lambda_{i_1}, \ldots, \lambda_{i_l}$ exhaust all elements from $\lambda_1, \ldots, \lambda_k$ that are not from $\mathit{dom}(D)$ and $\bar{b} \coloneqq \lambda_{j_1}, \ldots, \lambda_{j_m}$ exhaust those from $\lambda_1, \ldots, \lambda_k$ that are from $\mathit{dom}(D)$. Let $\varphi'(\bar{x}, y_{j_1}, \ldots, y_{j_m})$ be the formula $\exists y_{i_1}, \ldots, y_{i_l}\, (\alpha(\bar{x}, \bar{y}) \wedge \psi)$. Clearly, $\Sigma \models \forall \bar{x}, y_{j_1}, \ldots, y_{j_m}\, (\varphi'(\bar{x}, y_{j_1}, \ldots, y_{j_m}) \rightarrow \varphi(\bar{x}))$. Hence, we can create a successor of $(\bar{a}, \varphi(\bar{x}))$ that is labeled by $(\bar{a}, \bar{b}; \varphi'(\bar{x}, y_{j_1}, \ldots, y_{j_m}))$. Assume now that none of the $\lambda_1, \ldots, \lambda_k$ is from $\mathit{dom}(D)$. Furthermore, assume that $k \geq 1$, since otherwise we can just simply apply rule 2.(c) to reduce $q(\bar{x})$ to a conjunction of queries of the desired form. As in the proof of Lemma 44, there is a node v_0 in $\mathcal{F}$ whose label $\beta_0(\bar{a}, \bar{b})$ contains only values from $\mathit{dom}(D)$ as arguments and such that v_0 is an ancestor of the node labeling $\alpha_\mu(\bar{a}, \lambda_1, \ldots, \lambda_k)$. Furthermore, all the atoms from

$\mu(q(\bar{x}))$ that contain an element from $\lambda_1, \dots, \lambda_k$ as argument are also located in the subtree rooted at v_0. Let p be the query that results from deleting all atoms from $q(\bar{x})$ which are mapped via μ into the subtree rooted at v_0. Notice that p may be empty and has free variables among $\bar{x}$. Furthermore p is equivalent to a conjunction $p_1 \wedge \dots \wedge p_l$ of strictly acyclic queries. Let $\beta_0(\bar{x}, \bar{z})$ be the atomic formula whose image under an appropriate homomorphism equals $\beta_0(\bar{a}, \bar{b})$. A similar line of reasoning as in the proof of Lemma 44 shows that there are atomic formulas $\beta_1, \dots, \beta_m$ such that $var(\beta_i) \subseteq \{\bar{x}, \bar{z}\}$ and

$$\forall \bar{x}, \bar{z}\, (\beta_0(\bar{x}, \bar{z}) \wedge \beta_1 \wedge \dots \wedge \beta_m \wedge p \to \varphi(\bar{x})),$$

whence an application of rule 2.(b) and rule 2.(c) reduces the problem of constructing a derivation tree for $(\bar{a}, \varphi(\bar{x}))$ to that of constructing corresponding trees for $\beta_0, \dots, \beta_m$ and p. Notice that p is a conjunction of strictly guarded formulas of nesting depth at most n. Hence, the induction hypothesis guarantees the existence of such derivation trees. □

Having the above results in place, it is easy to show the following statement:

LEMMA 46. *Let D be a C-tree over $\mathbf{S}$ and $Q = (\mathbf{S}, \Sigma, q)$ an OMQ where Σ is guarded and q a BCQ. Then $D \models Q$ iff there is a squid decomposition $\delta = (q^+, \mu, H, A, V \coloneqq \{\bar{x}\})$ of q and a homomorphism $\eta\colon \mu(q^+) \to chase(D, \Sigma)$ such that:*

1. *$F \models H$ is witnessed by η, where F is the subinstance of $chase(D, \Sigma)$ induced by $dom(C)$.*
2. *There are strictly acyclic queries $q_1, \dots, q_l$ such that*
   (a) *$A^V(\bar{x}) \equiv q_1 \wedge \dots \wedge q_l$ and*
   (b) *for $i = 1, \dots, l$ and $\text{free}(q_i) = \{\bar{x}_i\}$, there are derivation trees for $(\eta(\bar{x}_i), q_i)$ with respect to D and Σ.*

PROOF. We can easily prove by induction on the number of chase steps that $chase(D, \Sigma)$ is an F-tree, where F is the subinstance of $chase(D, \Sigma)$ induced by $dom(C)$. Now the lemma at hand is immediate by combining this fact with Lemma 43 and Lemma 45. □

We are now ready to proceed with the proof of Lemma 24:

PROOF OF LEMMA 24. Lemma 46 will guide the construction of the 2WAPA we are now going to construct. Suppose $Q = (\mathbf{S}, \Sigma, q)$ is an OMQ from $(\mathbb{G}, \mathbb{BCQ})$ and let $l \geq 1$. We are going to construct a 2WAPA $\mathfrak{A}_{Q,l} = (S, \Gamma_{\mathbf{S},l}, \delta, s_0, \Omega)$ that accepts a consistent $\Gamma_{\mathbf{S},l}$-labeled tree t iff $[\![t]\!] \models Q$. In particular, the number of states of $\mathfrak{A}_{Q,l}$ will be at most exponential in the size of Q and at most polynomial in l, while the construction of $\mathfrak{A}_{Q,l}$ will be feasible in 2EXPTIME.

The state set. Let Λ denote the set of all Boolean acyclic queries over $\mathbf{S} \cup sch(\Sigma)$ that are of size at most $3|q|$. Notice that each of these queries is equivalent to a strictly guarded formula. Furthermore, assume that Λ is closed under V-reducts, for $V \subseteq var(q)$, provided that they are strictly acyclic as well, i.e., if $p \in \Lambda$ and $V \subseteq var(q)$, then also $p^V \in \Lambda$ provided p^V is strictly acyclic. For $\{\bar{a}\} \subseteq U_{\mathbf{S},l}$, let

$$\hat{S}(\bar{a}) \coloneqq \{p(\bar{x}/\bar{a}) \mid p \in \Lambda, \text{free}(p) = \{\bar{x}\}, |\bar{a}| = |\bar{x}|\}$$

and let $\hat{S}$ be the union of all the sets $\hat{S}(\bar{a})$. Now the set of states S consists of an initial state, denoted s_0, plus the set $\hat{S}$ factorized modulo logical equivalence. We denote by $[p]$ the equivalence class of a query $p \in \hat{S}$. Furthermore, for a strictly guarded formula φ, we may abuse notation and write $[\varphi]$ for the equivalence class of the strictly acyclic query $p \in \hat{S}$ that is equivalent to φ. Notice that the size of S is exponential in the size of Q, since there are only exponentially many CQs of size at most $3|q|$ that are mutually non-equivalent (cf. [10]).

The parity condition. We set $\Omega(s) \coloneqq 1$, for all $s \in S$. This means that only finite trees are accepted.

The transition function. In the following, for each $\rho \in \Gamma_{\mathbf{S},l}$, we denote by $\hat{\Theta}(\rho)$ the set of all pairs that are of the form $(\alpha_1 \wedge \dots \wedge \alpha_n, p_1 \wedge \dots \wedge p_m)$ for which there is a squid decomposition of the form $(q^+, \mu, H, T, \{\bar{x}\})$ and a function $\theta\colon \{\bar{x}\} \to \text{names}(\rho)$ such that:

- $H^{\{\bar{x}\}}(\theta(\bar{x})) \equiv \alpha_1 \wedge \dots \wedge \alpha_n$, where all the α_i are relational ground atoms.
- $T^{\{\bar{x}\}}(\theta(\bar{x})) \equiv p_1 \wedge \dots \wedge p_m$, where the p_i are strictly acyclic queries.

Call two pairs (φ_1, ψ_1) and (φ_2, ψ_2) as above *equivalent* if $\varphi_1 \equiv \varphi_2$ and $\psi_1 \equiv \psi_2$. Let $\Theta(\rho)$ be the set of equivalence classes under this relation and denote by $[\varphi, \psi]$ the equivalence of a pair (φ, ψ) under this relation. Now we fix for each $[p] \in S \setminus \{s_0\}$ a strictly guarded formula $\chi_{[p]}$ that is equivalent to all queries from $[p]$. Likewise, we fix a function $\vartheta_\rho\colon \Theta(\rho) \to \hat{\Theta}(\rho)$ such that $\vartheta_\rho([\varphi, \psi]) \in [\varphi, \psi]$, i.e., which picks a representative for each equivalence class $[\varphi, \psi]$.

Now let $\rho \in \Gamma_{\mathbf{S},l}$. Specify $\delta(\cdot, \rho)$ as follows:

1. For the initial state s_0, set

$$\delta(s_0, \rho) \coloneqq \bigvee \{\bigwedge_{i=1}^{n} [\alpha_i] \wedge \bigwedge_{i=1}^{m} [p_i] \mid (\alpha_1 \wedge \dots \wedge \alpha_n, p_1 \wedge \dots \wedge p_m) \in \vartheta_\rho(\Theta(\rho))\}.$$

   Intuitively, the automaton selects a squid decomposition where its components are instantiated by names occurring in the root node of the input tree. The automaton tries to verify the single compartments of the squid decomposition, i.e., it tries to match them to the chase expansion of the input database under Σ.

2. Let $[p] \in S \setminus \{s_0\}$. We define $\delta([p], \rho)$ according to a case distinction:

   (a) Suppose that $p \equiv \top$. Then $\delta([p], \rho) \coloneqq \mathsf{true}$.

   (b) Suppose $\chi_{[p]} = \exists \bar{y}\,(\alpha(\bar{a}, \bar{y}) \wedge \varphi)$, where $\alpha(\bar{a}, \bar{y})$ is an atomic formula (including equality), $\mathrm{free}(\varphi) \subseteq \{\bar{y}\}$, and $\bar{a}$ exhausts all names occuring in α. If $\{\bar{a}\} \not\subseteq \mathrm{names}(\rho)$ then $\delta([p], \rho) \coloneqq \mathsf{false}$. Otherwise,

$$\delta([p], \rho) \coloneqq \bigvee \{[\varphi(\bar{y}/\bar{b})] \mid \rho \models \alpha(\bar{a}, \bar{b}), \{\bar{b}\} \subseteq \mathrm{names}(\rho)\} \vee \Diamond[p] \;\vee \bigvee \mathrm{impl}(p, \rho),$$

   where

$$\begin{aligned} \mathrm{impl}(p, \rho) \coloneqq \{[p_1] \wedge \cdots \wedge [p_n] \mid\ & [p_1], \ldots, [p_n] \in S \setminus \{s_0\}, \{\bar{b}\} \subseteq \mathrm{names}(\rho), \\ & p_1 \wedge \cdots \wedge p_n \equiv q, \\ & \Sigma \models \forall \bar{x}, \bar{y}\,(q(\bar{a}/\bar{x}, \bar{b}/\bar{y}) \rightarrow p(\bar{a}/\bar{x}))\}. \end{aligned}$$

   We provide some intuitive explanation for this second case.

   (a) If p is the empty query, it can be satisfied at any input node and, hence, the automaton accepts unconditionally on this computation branch.

   (b) Otherwise, we first inspect the strictly guarded formula $\chi_{[p]}$ at hand. If the names occurring in the guard $\alpha(\bar{a}, \bar{y})$ are not present at the current node, it rejects. Otherwise, it tries to satisfy $\alpha(\bar{a}, \bar{y})$ with all possible assignments for $\bar{y}$ at the current node and then proceed in state $[\varphi(\bar{y}/\bar{b})]$. Apart from these possibilities, the automaton can decide to move to any neighboring node (i.e., the parent or a child) while remaining in state $[p]$. This amounts to an exhaustive search of the input tree that tries to satisfy p in the input tree. Furthermore, the automaton may choose to construct derivation trees for p. There, it uses the information provided by Σ in order to find strictly acyclic queries $p_1, \ldots, p_n$ that imply p. Consequently, it tries to proceed its search with $[p_1], \ldots, [p_n]$.

We shall now briefly comment on the running time needed to construct $\mathfrak{A}_{Q,l}$. The interesting part of the construction concerns the transition function δ, in particular point 2.(b) involving $\mathrm{impl}(p, \rho)$. We have seen that in the proofs of Lemma 44 and Lemma 45 that there are double-exponentially many candidates for the query $q(\bar{a}/\bar{x}, \bar{b}/\bar{y})$ that (possibly) implies $p(\bar{a}/\bar{x})$ under Σ. Furthermore, $q(\bar{a}/\bar{x}, \bar{b}/\bar{y})$ consists of at most exponentially many atoms. Each check whether such a query q at hand implies p requires at most double-exponential time in the size of p. This follows from the well-known fact that checking query implication under a set of guarded rules is feasible in 2ExpTime with respect to the size of the right-hand side query, and in polynomial time with respect to the size of the left-hand side query (cf. [23]), i.e., the *data complexity* of query answering under guarded tgds is polynomial time. □

PROOFS OF SECTION 6

Proof of Theorem 26

A proof sketch is given in the main body of the paper. However, the fact that $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}), (\mathbb{S}, \mathbb{CQ}))$ is in 2ExpTime deserves a formal proof. Recall that to establish the latter result we need a more refined complexity analysis of the problem of deciding whether a guarded OMQ is contained in a UCQ; this is discussed in the main body of the paper. In fact, it suffices to show the following result. As in the previous section, we focus on constant-free tgds and CQs, but all the results can be extended to the general case at the price of more involved definitions and proofs. Moreover, we assume that tgds have only one atom in the head. Recall that we write $var_{\geq 2}(q)$ for the variables of q that appear in more than one atom, and we also write $var_{=1}(q)$ for the variables of q that appear only in one atom. Then:

Proposition 47. *Consider $Q \in (\mathbb{G}, \mathbb{BCQ})$ and a Boolean CQ q. The problem of deciding whether $Q \subseteq q$ is feasible in*

1. *double-exponential time in $(||Q|| + |var_{\geq 2}(q)|)$; and*
2. *exponential time in $|var_{=1}(q)|$.*

It is easy to verify that the above result, together with the algorithm devised in the main body of the paper, implies that $\mathsf{Cont}((\mathbb{G}, \mathbb{CQ}), (\mathbb{S}, \mathbb{CQ}))$ is in 2ExpTime. The rest of this section is devoted to show the above proposition. Our crucial task is, given a CQ q, to devise an automaton that accepts consistent labeled trees which correspond to databases that make q true.

Lemma 48. *Let q be a Boolean CQ over $\mathbf{S}$. There is a 2WAPA $\mathfrak{A}_{q,l}$, where $l > 0$, that accepts a consistent $\Gamma_{\mathbf{S},l}$-labeled tree t iff $[\![t]\!] \models q$. The number of states of $\mathfrak{A}_{q,l}$ is exponential in $|var_{\geq 2}(q)|$ and polynomial in $(|var_{=1}(q)| + ar(\mathbf{S}) + l)$. Furthermore, $\mathfrak{A}_{q,l}$ can be constructed in exponential time.*

Proof. We are going to construct $\mathfrak{A}_{q,l} = (S, \Gamma_{\mathbf{S},l}, \delta, s_0, \Omega)$. Let $x_1, \ldots, x_n$ be the variables of $var_{=1}(q)$ and fix a total order $x_1 \prec x_2 \prec \cdots \prec x_n$ among them. Define the state set S to be

$$S \coloneqq \{s_{y,\theta} \mid \theta \colon V \rightarrow U_{\mathbf{S},l}, V \subseteq var_{\geq 2}(q), y \in var_{=1}(q) \cup \{\sharp\}\}.$$

Notice that $|S| = O(|var_{=1}(q)| \cdot (ar(\mathbf{S}) + l)^{|var_{\geq 2}(q)|})$. We set $s_0 \coloneqq s_{\sharp,\varnothing}$, where $\varnothing$ denotes the empty substitution. In the following, we treat q as a set of relational atoms and let $X = var_{\geq 2}(q)$. For $\rho \in \Gamma_{\mathbf{S},l}$ and $s_{y,\theta} \in S$, define $\delta(s_{y,\theta}, \rho)$ as follows:

- If $y = \sharp$, distinguish the following cases:
  1. If there is an atom $\alpha \in \theta(q)$ such that $var(\alpha) \cap X \neq \varnothing$ and $dom(\alpha) \cap U_{\mathbf{S},l} \not\subseteq \text{names}(\rho)$, then $\delta(s_{\sharp,\theta}, \rho) \coloneqq \mathsf{false}$.
  2. Otherwise, let

$$\delta(s_{\sharp,\theta}, \rho) \coloneqq \begin{cases} \bigvee\{s_{\sharp,\eta} \mid \eta \supseteq \theta, \rho \models \exists \bar{x} \bigwedge(\eta(q) \setminus \theta(q))\} \vee \Diamond s_{\sharp,\theta}, & \text{if } X \cap var(\theta(q)) \neq \varnothing, \\ s_{x_1,\theta}, & \text{otherwise.} \end{cases}$$

- Suppose $y = x_i$, for some $i = 1, \ldots, n$. Let $\alpha_{i,\theta}$ denote the unique atom $\alpha \in \theta(q)$ such that $x_i \in var(\alpha)$. Set

$$\delta(s_{x_i,\theta}, \rho) \coloneqq \begin{cases} s_{x_{i+1},\theta}, & \text{if } \rho \models \exists \bar{x}\, \alpha_{i,\theta} \text{ and } i < n, \\ \mathsf{true}, & \text{if } \rho \models \exists \bar{x}\, \alpha_{i,\theta} \text{ and } i = n, \\ \Diamond s_{x_i,\theta}, & \text{otherwise.} \end{cases}$$

Set the parity condition Ω to be $\Omega(s) \coloneqq 1$ for all $s \in S$. Intuitively, the automaton works in two passes. The first pass consists of the runs working on states of the form $s_{\sharp,\theta}$. In this pass, the automaton tries to find an assignment for the variables in the query that appear in at least two distinct atoms. When a candidate assignment θ is found, the automaton changes to state $s_{x_1,\theta}$ which is the beginning of the second pass. A state of the form $s_{x_i,\theta}$ means that the assignment θ can be extended to all variables $x \prec x_i$ and, in this state, the automaton tries to extend θ to cover the variable x_i. The automaton accepts if it is able to extend the candidate assignment θ to all $x_1, \ldots, x_n$. □

Having the above result in place, we can now reduce the problem in question to the emptiness problem for 2WAPA.

LEMMA 49. *Consider $Q \in (\mathbb{G}, \mathbb{BCQ})$ and a Boolean CQ q. We can construct in double-exponential time in $||Q||$ and in exponential time in $||q||$ a 2WAPA $\mathfrak{A}$, which has exponentially many states in $(||Q|| + |var_{\geq 2}(q)|)$ and polynomially many states in $|var_{=1}(q)|$, such that*

$$Q \subseteq q \iff \mathcal{L}(\mathfrak{A}) = \varnothing.$$

PROOF. Let $Q = (\mathbf{S}, \Sigma, q')$ and $l = ar(\mathbf{S} \cup sch(\Sigma)) \cdot |q'|$. Then $\mathfrak{A}$ is defined as:

$$(\mathfrak{C}_{\mathbf{S},l} \cap \mathfrak{A}_{Q,l}) \cap \overline{\mathfrak{A}_{q,l}}.$$

It is an easy task to verify that the claim follows from Lemmas 22, 23, 24, and 48. □

It is clear that Proposition 47 is an easy consequence of Lemma 49.

PROOFS OF SECTION 7

Recall that we focus on unary and binary predicates. Moreover, we consider constant-free tgds and CQs, and we assume that tgds have only one atom in the head.

Proof of Proposition 30

Basics. Let D be a C-tree of width two. We say that a tree decomposition $\delta = (\mathcal{T}, (X_t)_{t \in T})$ witnessing that D is a C-tree is *lean*, if it satisfies the following conditions:

- The elements from $dom(C)$ occur only in the root of $\mathcal{T}$ and its immediate successors.
- If w is a child of v in $\mathcal{T}$, then there are unique $c, d \in dom(D)$ such that $X_v \cap X_w = \{c\}$ and $X_w \setminus X_v = \{d\}$. The element d is called *new at* w.
- It follows from the previous item that every node $v \neq \varepsilon$ in $\mathcal{T}$ has a unique new element $c \in dom(D)$. We additionally require that c appears in the bag of each child of v.

Intuitively, C-trees D that have lean tree decomposition represent the *actual* tree structure of D. It is fairly straightforward to see that every C-tree has a lean tree decomposition.

Recall that the *Gaifman graph* of D is the graph $\mathcal{G}(D) = (V, E)$ with $V \coloneqq dom(D)$ and $(a, b) \in E$ if a and b coexist in some atom of D. Given two nodes a, b from $\mathcal{G}(D)$, the *distance from a to b* in $\mathcal{G}(D)$, denoted $d_{\mathcal{G}(D)}(a, b)$, is the minimum length of a path between a and b, and ∞ if such a path does not exist. For $a, b \in dom(D)$, we denote by $d_\delta(a, b)$ the minimum distance among two nodes of $\mathcal{T}$ that respectively have a and b in their bags. We call $d_\delta(a, b)$ the *distance from a to b in δ*.

Notice that in a tree decomposition δ witnessing that D is a C-tree, any element $a \in dom(D)$, if a appears in the bag of v, then it occurs only at v, at v and its children, or at v and its parent. Since furthermore the bag of the root node is uniquely determined by C, each node in the tree has a uniquely determined set of child nodes whose bags are determined by the structure of D alone. Therefore, the following two lemmas follow immediately.

LEMMA 50. *Let $\delta = (\mathcal{T}, (X_t)_{t \in T})$ be a lean tree decomposition witnessing that D is a C-tree. Then $d_\delta(a, b) \leq d_{\mathcal{G}(D)}(a, b)$ for all $a, b \in dom(D)$.*

LEMMA 51. *Let δ and δ' be two lean tree decompositions witnessing that D is a C-tree. Then $d_\delta(a,b) = d_{\delta'}(a,b)$ for all $a, b \in dom(D)$.*

In the following, we denote by $D_{\leq k}$ the subinstance of D induced by the set of elements whose distance from any $a \in dom(C)$ in *any* lean tree decomposition δ is bounded by k. The subinstance $D_{>k}$ is defined analogously.The *branching degree* of a lean tree decomposition is the maximum number of child nodes of any node contained in the tree of δ. Notice that two lean tree decompositions of a C-tree D always have the same branching degree; the argument is similar as for the two lemmas above. Hence, we can simply speak about *the* branching degree of D.

Encodings. Recall that a consistent $\Gamma_{\mathbf{S},l}$-labeled tree $t = (T, \mu)$ encodes information on an $\mathbf{S}$-database D *and* an appropriate tree decomposition δ of D. It is clear that $[\![t]\!]$ has a lean tree decomposition, but it is not guaranteed that this is reflected in δ as well. We call (the consistent) t *lean*, if the tree decomposition $\delta_t = (\mathcal{T} \coloneqq (T, E), (X_v)_{v \in T})$ is, where xEy iff $y = x \cdot i$ for some $i \in \mathbb{N} \setminus \{0\}$ and $X_v \coloneqq \{[v]_a \mid a \in \text{names}(v)\}$. The following is easy to prove:

LEMMA 52. *There is a 2WAPA on trees $\mathfrak{L}_{\mathbf{S},l}$ that accepts a consistent $\Gamma_{\mathbf{S},l}$-labeled tree iff it is lean. The number of states of $\mathfrak{L}_{\mathbf{S},l}$ is bounded logarithmically in the size of $\Gamma_{\mathbf{S},l}$ and $\mathfrak{L}_{\mathbf{S},l}$ can be constructed in polynomial time in the size of $\Gamma_{\mathbf{S},l}$.*

Let $t = (T, \mu)$ be a labeled tree. The *branching degree* of a node $x \in T$ is the cardinality of $\{i \mid x \cdot i \in T, i \in \mathbb{N} \setminus \{0\}\}$; the branching degree of t is the maximum over all branching degrees of its nodes and ∞ is this maximum does not exist. We also say that t is *m-ary* if the branching degree of t is bounded by m. A node $x \in T$ is a *leaf node* of t if it has branching degree zero. The *depth* of T is the maximum length among the lengths of all branches and ∞ if this maximum does not exist. Let us remark that the branching degree of the lean $\Gamma_{\mathbf{S},l}$-labeled tree t as defined for labeled trees equals the branching degree of $[\![t]\!]$ as defined above.

LEMMA 53. *Let $Q = (\mathbf{S}, \Sigma, q)$ be an OMQ from $(\mathbb{G}_2, \mathbb{BCQ})$. There is an $m \geq 0$ such that the following are equivalent:*

1. *There is an $\mathbf{S}$-database D such that $D \models Q$.*
2. *There is a C-tree $\hat{D}$ with $|dom(C)| \leq 2|q|$ and branching degree at most m such that $\hat{D} \models Q$.*

PROOF. Let $l \coloneqq 2|q|$ and, let $\mathfrak{A}_{Q,l}$ be the 2WAPA from Lemma 24. Take the intersection of $\mathfrak{A}_{Q,l}$ with (i) the 2WAPA $\mathfrak{C}_{\mathbf{S},l}$ from Lemma 23 and (ii) the 2WAPA from Lemma 52 that checks leanness. Call the resulting automaton $\mathfrak{B}$. Then $\mathfrak{B}$ accepts a $\Gamma_{\mathbf{S},l}$-labeled tree t iff t is lean and consistent and $[\![t]\!] \models Q$. We let m be the number of states of $\mathfrak{B}$ and claim that this is the required bound on the branching degree.

First of all, notice that the first item of the lemma trivially implies the second independently from the choice of m. For the other direction, suppose that $D \models Q$ for some $\mathbf{S}$-database D. Then there is a C-tree B such that $dom(C) \leq 2|q|$ and $B \models Q$. Being a C-tree, B has a lean tree decomposition δ and the encoding of B together with δ corresponds to a lean $\Gamma_{\mathbf{S},l}$-labeled tree t. It follows that $t \in \mathcal{L}(\mathfrak{B})$. By the results of [44], it follows that there is a $t' \in \mathcal{L}(\mathfrak{B})$ whose branching degree is bounded by the number of states of $\mathfrak{B}$, i.e., by m. The tree t' is lean and consistent, therefore $[\![t']\!]$ is a C'-tree of branching degree at most m for some $C' \subseteq [\![t]\!]$ such that $|dom(C')| \leq 2|q|$. Furthermore, $[\![t']\!] \models Q$, as required. □

We are now ready to prove Proposition 30:

PROOF OF PROPOSITION 30. We largely follow [16] here. Choose m as in Lemma 53 above. Suppose first that Q is UCQ rewritable. Let $p \coloneqq p_1 \vee \cdots \vee p_n$ be a corresponding UCQ rewriting. Since the query q is connected, we can assume that p is as well. We choose $k > \max\{|p_i| : i = 1, \ldots, n\}$ and suppose that $D \models Q$ for some C-tree D. Since p is a UCQ rewriting, $D \models p_i$ for some $i = 1, \ldots, n$. Fix a homomorphism μ witnessing that $D \models p_i$. We distinguish cases. Suppose first that $\mu(var(p_i)) \cap dom(C) \neq \varnothing$. Since p is connected, it follows $D_{\leq k} \models p_i$ by Lemma 50 and so $D_{\leq k} \models p$. On the other hand, if $\mu(var(p)) \cap dom(C) = \varnothing$, then it is also easy to check that $D_{>0} \models p$.

For the other direction, suppose that the second item of the proposition's statement holds, i.e., there is a $k \geq 0$ such that for all C-trees D over $\mathbf{S}$ with $|dom(C)| \leq 2|q|$ and branching degree at most m it holds that $D \models Q$ implies $D_{\leq k} \models Q$ or $D_{>0} \models Q$. Let Λ be the set of all C-trees such that $|dom(C)| \leq 2|q|$ and that have branching degree at most m such that $D \models Q$. We regard Λ as a set of BCQs and regard it as factorized modulo logical equivalence. It is clear that Λ is finite then and we claim that $p \coloneqq \bigvee_{i=1}^{n} p_i$ is a UCQ rewriting of Q. We explicitly include the case where Λ is empty, in which case p is equivalent to the empty disjunction $\bot$ and there is no database D at all such that $D \models Q$. To see that p is indeed a UCQ rewriting of Q, let D be an $\mathbf{S}$-database such that $D \models p$. Then there is an $i = 1, \ldots, n$ such that $D \models p_i$. Furthermore, $[p_i] \models Q$ and so $D \models Q$ as well, since Q is closed under homomorphisms. Suppose now $D \models Q$. We know that there is a C-tree $\hat{D}$ with $|dom(C)| \leq l \coloneqq 2|q|$ and branching degree at most m such that $\hat{D} \models Q$ and—when we regard $\hat{D}$ as an instance—there is a homomorphism from $\hat{D}$ to D. Let $D' \subseteq \hat{D}$ be a minimal connected subset of $\hat{D}$ such that $D' \models Q$. D' is again a C'-tree for some $C' \subseteq D'$. Therefore $D'_{\leq k} \models Q$ or $D'_{>0} \models Q$. The latter is impossible by minimality of D'. Hence, $D'_{\leq k} \models Q$ and so there is a (logically equivalent) copy of $D'_{\leq k}$ contained in Λ. Hence, $D'_{\leq k} \models p$, therefore $D' \models p$, and hence $\hat{D} \models p$. Recall that, when $\hat{D}$ is regarded as an instance, there is a homomorphism from $\hat{D}$ to D. Therefore, $D \models p$. □

Proof of Proposition 31

Let $Q = (\mathbf{S}, \Sigma, q)$ be an OMQ from $(\mathbb{G}_2, \mathbb{BCQ})$ such that q is connected. We are going to show that the desired 2WAPA $\mathfrak{A}$ can be constructed in 2ExpTime. Notice that, using similar results as in [16], this gives us a decision procedure for deciding $\mathsf{UCQRew}(\mathbb{G}_2, \mathbb{CQ})$ also for non-connected queries. Let us first introduce some auxiliary notions.

2WAPAs on m-ary trees. A 2WAPA $\mathfrak{B}$ *on m-ary* trees is just defined as a 2WAPA, except that its transitions $\mathsf{tran}(\mathfrak{B})$ are $\{\langle k\rangle s, [k]s \mid -1 \leq k \leq m, s \in S\}$, where S is the state set of $\mathfrak{B}$. The notion of run is then defined on m-ary trees only and its definition is modified in the obvious way so as to deal with the transitions $\langle k\rangle s, [k]s$. Intuitively, for $k = 1, \dots, m$, a transition $\langle k\rangle s$ means that the automaton should move to the k-th child of the current node (which is then required to exist) and assume state s. Correspondingly, $[k]s$ means that the automaton should move to the k-th child and assume state s provided that this k-th child exists at all. We remark that all 2WAPAs constructed in this paper so far can easily be modified to work on m-ary trees as well and we shall assume in the following that they do so. Furthermore, deciding whether $\mathcal{L}(\mathfrak{B})$ is feasible in exponential time in the number of states of $\mathfrak{B}$ and in polynomial time in the size of the input alphabet of $\mathfrak{B}$ (cf. [57]).

Let m be as in Proposition 30. In the following, we shall regard all trees mentioned in the following as m-ary and let $l \coloneqq 2|q|$. Before proceeding to a proof of Proposition 31, we must make the notion of being an "extension" of a labeled tree more precise.

Extensions of trees. Let $\mathfrak{B}_Q$ be a 2WAPA that accepts a $\Gamma_{\mathbf{S},l}$-labeled tree t iff (i) t is lean and consistent, (ii) $[\![t]\!] \models Q$, and (iii) $[\![t]\!]_{>0} \not\models Q$. Notice that a 2WAPA $\mathfrak{A}_Q^{>0}$ that accepts a lean and consistent $\Gamma_{\mathbf{S},l}$-labeled tree iff $[\![t]\!]_{>0} \not\models Q$ can be easily constructed using the construction in Lemma 24. Hence, $\mathfrak{B}_Q$ can be constructed intersecting several 2WAPAs we have already encountered.

Let Π be the set of all tuples of the form (s, s'), where s and s' are states of $\mathfrak{B}_Q$. We define a new alphabet $\Lambda \coloneqq 2^{\mathbb{K}_{\mathbf{S},l} \cup \Pi}$. Notice that Λ is of double-exponential size in the size of Q. For $\rho \in \Lambda$, we denote by $\rho \restriction \Gamma_{\mathbf{S},l}$ the *restriction of ρ to* $\Gamma_{\mathbf{S},l}$, that is, $\rho \cap \mathbb{K}_{\mathbf{S},l}$. The *restriction* of a Λ-labeled tree t to $\Gamma_{\mathbf{S},l}$, denoted $t \restriction \Gamma_{\mathbf{S},l}$, is the tree that arises from t when we restrict the label of each node of t to $\Gamma_{\mathbf{S},l}$. We say that a Λ-labeled tree is *consistent* if (i) its restriction to $\Gamma_{\mathbf{S},l}$ is consistent and (ii) symbols $\rho \in \Lambda$ such that $\rho \cap \Pi \neq \varnothing$ appear only in leaf nodes of t. Likewise, we say that a consistent t is *lean* if $t \restriction \Gamma_{\mathbf{S},l}$ is. The decoding $[\![t]\!]$ of t is naturally extended to consistent Λ-labeled trees by setting $[\![t]\!] \coloneqq [\![t \restriction \Gamma_{\mathbf{S},l}]\!]$. The following lemma is a straightforward extension of Lemmas 23 and 52.

Lemma 54. *There are 2WAPAs $\mathfrak{C}_\Lambda$ and $\mathfrak{L}_\Lambda$ that respectively accept a Λ-labeled tree iff it is consistent and lean. Both have logarithmically many states in the size of Λ and can be constructed in polynomial time in the size of Λ.*

Let t be a lean and consistent Λ-labeled tree. We say that t' is an *extension* of t if t' is a $\Gamma_{\mathbf{S},l}$-labeled tree that arises from t by attaching $\Gamma_{\mathbf{S},l}$-labeled trees to those leaves of t that contain elements from Π. Furthermore, for such nodes, the labels of the corresponding nodes in t' are those of t restricted to $\Gamma_{\mathbf{S},l}$.

Definition 14. Let $\mathcal{L}_Q$ be the set of all lean and consistent Λ-labeled trees t such that $[\![t]\!] \not\models Q$, yet there is an extension t' of t such that $[\![t']\!] \models Q$ and $[\![t']\!]_{>0} \not\models Q$. ∎

Lemma 55. *$\mathcal{L}_Q$ is infinite iff Q is not UCQ rewritable.*

Proof. Suppose $\mathcal{L}_Q$ is infinite. Since the trees at hand have bounded branching degree, for every $k \geq 0$, there is a $t \in \mathcal{L}_Q$ such that $[\![t]\!]$ is a C-tree (for some $C \subseteq [\![t]\!]$) that contains individuals whose distance from any $a \in dom(C)$ is greater than or equal to k and $[\![t]\!] \not\models Q$, yet for some extension t' of t, we have $[\![t']\!] \models Q$ but $[\![t']\!]_{>0} \not\models Q$. Suppose now that Q is UCQ rewritable. Let ℓ be such that for all C'-trees D (of the appropriate dimensions), $D \models Q$ implies $D_{\leq \ell} \models Q$ or $D_{>0} \models Q$. Choose $k > \ell$ and t, t' such that (i) t' is an extension of t, (ii) t has depth greater than k, and (iii) $[\![t]\!] \not\models Q$ but $[\![t']\!] \models Q$ and $[\![t']\!]_{>0} \not\models Q$. Since $[\![t']\!] \models Q$, we know that $[\![t']\!]_{\leq \ell} \models Q$ or $[\![t']\!]_{>0} \models Q$. The latter is impossible by assumption, the former contradicts the fact $[\![t]\!] \not\models Q$, since $k > \ell$. This proves the direction from left to right. The other direction is immediate. □

We are now ready to establish Proposition 31:

Proof of Proposition 31. We are now going to describe the construction of a 2WAPA $\mathfrak{A}$ such that $\mathcal{L}(\mathfrak{A}) = \mathcal{L}_Q$, which will prove the claim by virtue of Lemma 55. This automaton is the intersection of several ones. First of all, we ensure that all the accepted Λ-trees are lean and consistent (cf. Lemma 54). We additionally intersect the automaton with the complement of $\mathfrak{A}_{Q,l}$ from Lemma 23 (more precisely, the version of it running on Λ-labeled trees) and another automaton $\mathfrak{D}_Q = (S, \Lambda, \delta, s_0, \Omega)$ whose construction we shall describe in more detail here. On a high level, $\mathfrak{D}_Q$ will be constructed so as to accept a lean and consistent Λ-labeled tree if and only if there is an extension t' of t such that $\mathfrak{B}_Q$ accepts t'. Let $\hat{S}$ be the set of states of $\mathfrak{B}_Q$, $\hat{\delta}$ its transition function, and $\hat{\Omega}$ its parity function. For $\sigma \in \Gamma_{\mathbf{S},l}$, let $\hat{B}(\sigma)$ be the set of tuples $(s, s') \in \hat{S} \times (\hat{S} \cup \{\mathsf{true}\})$ such that the following holds:

- There is a $\Gamma_{\mathbf{S},l}$-labeled tree $t = (T, \eta)$ such that $\eta(\varepsilon) = \sigma$ and a run $t_r = (T_r, \eta_r)$ of $\mathfrak{B}_Q$ on t such that
  1. $\eta_r(\varepsilon) = (\varepsilon, s)$, i.e., t_r *starts from* s;[12]

[12] Strictly speaking, t_r is, of course, not a run since it does not start in the initial state.

2. $s' = \mathsf{true}$ and t_r is accepting on $\mathfrak{B}_Q$, or there is a node $v \in T_r$ such that $\eta_r(v) = (\varepsilon, s')$.

Now the set of states of $\mathfrak{D}_Q$ is the same as of $\mathfrak{B}_Q$, i.e., $S \coloneqq \hat{S}$. Accordingly, the initial state of $\mathfrak{D}_Q$ is that of $\mathfrak{B}_Q$. Furthermore $\Omega(s) \coloneqq \hat{\Omega}(s)$, for every $s \in S$. Given $s \in S$ and $\rho \in \Lambda$, we let

$$\delta(s, \rho) \coloneqq \bigvee \{s' \mid (s, s') \in \rho \cap \hat{B}(\Gamma_{\mathbf{S},l} \restriction \rho)\} \vee \hat{\delta}(s, \Gamma_{\mathbf{S},l} \restriction \rho).$$

We are going to give an intuitive explanation of this construction in the following. Roughly, a pair $(s, s') \in \hat{B}(\sigma)$ indicates that there is a $\Gamma_{\mathbf{S},l}$-labeled tree t and run of $\mathfrak{B}_Q$ on t such that the root of t is labeled with σ, the run starts in state s, and either $\mathfrak{B}_Q$ accepts t, or it traverses the root again at some point, then being in state s'. The set $\hat{B}(\sigma)$ can be computed a priori in 2ExpTime; considering that $\Gamma_{\mathbf{S},l}$ is of double-exponential size in the size of Q, it follows that the collection $\{\hat{B}(\sigma)\}_{\sigma \in \Gamma_{\mathbf{S},l}}$ can be computed in 2ExpTime. Now the input tree for $\mathfrak{D}_Q$ comes with labels from Π of the form (s, s') in its leaves. These "types" amount to guesses of possible extensions of the input tree. Utilizing the sets $\hat{B}(\sigma)$, $\mathfrak{D}_Q$ thus explores the possible ways how the given input tree can be extended to a $\Gamma_{\mathbf{S},l}$-labeled tree t' that is accepted by $\mathfrak{B}_Q$. □

9. REFERENCES

[1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.

[2] A. Amarilli, M. Benedikt, P. Bourhis, and M. Vanden Boom. Query answering with transitive and linear-ordered data. In *IJCAI*, pages 893–899, 2016.

[3] T. J. Ameloot, B. Ketsman, F. Neven, and D. Zinn. Weaker forms of monotonicity for declarative networking: A more fine-grained answer to the calm-conjecture. In *PODS*, pages 64–75, 2014.

[4] T. J. Ameloot, B. Ketsman, F. Neven, and D. Zinn. Datalog queries distributing over components. In *ICDT*, pages 308–323, 2015.

[5] T. J. Ameloot, F. Neven, and J. V. den Bussche. Relational transducers for declarative networking. *J. ACM*, 60(2):15, 2013.

[6] M. Arenas, R. Hull, W. Martens, T. Milo, and T. Schwentick. Foundations of Data Management (Dagstuhl perspectives workshop 16151). *Dagstuhl Reports*, 6(4):39–56, 2016.

[7] F. Baader, M. Bienvenu, C. Lutz, and F. Wolter. Query and predicate emptiness in ontology-based data access. *J. Artif. Intell. Res. (JAIR)*, 56:1–59, 2016.

[8] J.-F. Baget, M. Leclère, M.-L. Mugnier, and E. Salvat. On rules with existential variables: Walking the decidability line. *Artif. Intell.*, 175(9-10):1620–1654, 2011.

[9] V. Bárány, M. Benedikt, and B. ten Cate. Rewriting guarded negation queries. In *MFCS*, pages 98–110, 2013.

[10] V. Bárány, G. Gottlob, and M. Otto. Querying the guarded fragment. *Logical Methods in Computer Science*, 10(2), 2014.

[11] P. Barceló, G. Berger, and A. Pieris. Containment for rule-based ontology-mediated queries. *CoRR*, abs/1703.07994, 2017.

[12] C. Beeri and M. Y. Vardi. The implication problem for data dependencies. In *ICALP*, pages 73–85, 1981.

[13] M. Benedikt, P. Bourhis, and M. Vanden Boom. A step up in expressiveness of decidable fixpoint logics. In *LICS*, pages 817–826, 2016.

[14] M. Benedikt and G. Gottlob. The impact of virtual views on containment. *PVLDB*, 3(1):297–308, 2010.

[15] G. Berger and A. Pieris. Ontology-mediated queries distributing over components. In *IJCAI*, pages 943–949, 2016.

[16] M. Bienvenu, P. Hansen, C. Lutz, and F. Wolter. First order-rewritability and containment of conjunctive queries in horn description logics. In *IJCAI*, pages 965–971, 2016.

[17] M. Bienvenu, C. Lutz, and F. Wolter. Query containment in description logics reconsidered. In *KR*, 2012.

[18] M. Bienvenu, C. Lutz, and F. Wolter. First-order rewritability of atomic queries in horn description logics. In *IJCAI*, pages 754–760, 2013.

[19] M. Bienvenu, B. ten Cate, C. Lutz, and F. Wolter. Ontology-based data access: a study through disjunctive datalog, csp, and MMSNP. In *PODS*, pages 213–224, 2013.

[20] P. Bourhis, M. Krötzsch, and S. Rudolph. Reasonable highly expressive query languages. In *IJCAI*, pages 2826–2832, 2015.

[21] P. Bourhis, M. Manna, M. Morak, and A. Pieris. Guarded-based disjunctive tuple-generating dependencies. *ACM Trans. Database Syst.*, 41(4), 2016.

[22] P. Buneman, S. B. Davidson, G. G. Hillebrand, and D. Suciu. A query language and optimization techniques for unstructured data. In *SIGMOD.*, pages 505–516, 1996.

[23] A. Calì, G. Gottlob, and M. Kifer. Taming the infinite chase: Query answering under expressive relational constraints. In *KR*, pages 70–80, 2008.

[24] A. Calì, G. Gottlob, and M. Kifer. Taming the infinite chase: Query answering under expressive relational constraints. *J. Artif. Intell. Res.*, 48:115–174, 2013.

[25] A. Calì, G. Gottlob, and T. Lukasiewicz. A general Datalog-based framework for tractable query answering over ontologies. *J. Web Sem.*, 14:57–83, 2012.

[26] A. Calì, G. Gottlob, and A. Pieris. Advanced processing for ontological queries. *PVLDB*, 3(1):554–565, 2010.
[27] A. Calì, G. Gottlob, and A. Pieris. Towards more expressive ontology languages: The query answering problem. *Artif. Intell.*, 193:87–128, 2012.
[28] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *J. Autom. Reasoning*, 39(3):385–429, 2007.
[29] A. K. Chandra and P. M. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *STOC*, pages 77–90, 1977.
[30] S. Chaudhuri, R. Krishnamurthy, S. Potamianos, and K. Shim. Optimizing queries with materialized views. In *ICDE*, pages 190–200, 1995.
[31] S. Chaudhuri and M. Y. Vardi. On the equivalence of recursive and nonrecursive datalog programs. *J. Comput. Syst. Sci.*, 54(1):61–78, 1997.
[32] S. S. Cosmadakis, H. Gaifman, P. C. Kanellakis, and M. Y. Vardi. Decidable optimization problems for database logic programs (preliminary report). In *STOC*, pages 477–490, 1988.
[33] E. Dantsin and A. Voronkov. Complexity of query answering in logic databases with complex values. In *LFCS*, pages 56–66, 1997.
[34] T. Eiter, T. Lukasiewicz, and L. Predoiu. Generalized consistent query answering under existential rules. In *KR*, pages 359–368, 2016.
[35] R. Fagin, P. G. Kolaitis, R. J. Miller, and L. Popa. Data exchange: Semantics and query answering. *Theor. Comput. Sci.*, 336(1):89–124, 2005.
[36] R. Fagin and M. Y. Vardi. The theory of data dependencies - an overview. In *ICALP*, pages 1–22, 1984.
[37] M. F. Fernández, D. Florescu, A. Y. Levy, and D. Suciu. Verifying integrity constraints on web sites. In *IJCAI*, pages 614–619, 1999.
[38] J. Flum, M. Frick, and M. Grohe. Query evaluation via tree-decompositions. *J. ACM*, 49(6):716–752, 2002.
[39] M. Friedman, A. Y. Levy, and T. D. Millstein. Navigational plans for data integration. In *AAAI/IAAI*, pages 67–73, 1999.
[40] G. Gottlob, G. Orsi, and A. Pieris. Query rewriting and optimization for ontological databases. *ACM Trans. Database Syst.*, 39(3):25:1–25:46, 2014.
[41] G. Gottlob and C. H. Papadimitriou. On the complexity of single-rule datalog queries. *Inf. Comput.*, 183(1):104–122, 2003.
[42] G. Gottlob, A. Pieris, and M. Simkus. The impact of active domain predicates on guarded existential rules. In *RR*, pages 94–110, 2016.
[43] G. Gottlob, S. Rudolph, and M. Simkus. Expressiveness of guarded existential rule languages. In *PODS*, pages 27–38, 2014.
[44] E. Grädel and I. Walukiewicz. Guarded fixed point logic. In *LICS*, pages 45–54, 1999.
[45] T. Imielinski and W. L. Jr. Incomplete information in relational databases. *J. ACM*, 31(4):761–791, 1984.
[46] D. S. Johnson. A catalog of complexity classes. In *Handbook of Theoretical Computer Science*, pages 67–161. 1990.
[47] D. S. Johnson and A. C. Klug. Testing containment of conjunctive queries under functional and inclusion dependencies. *J. Comput. Syst. Sci.*, 28(1):167–189, 1984.
[48] T. Lukasiewicz, M. V. Martinez, A. Pieris, and G. I. Simari. From classical to consistent query answering under existential rules. In *AAAI*, pages 1546–1552, 2015.
[49] D. Maier, A. O. Mendelzon, and Y. Sagiv. Testing implications of data dependencies. *ACM Trans. Database Syst.*, 4(4):455–469, 1979.
[50] S. Malik and L. Zhang. Boolean satisfiability: from theoretical hardness to practical success. *CACM*, 52(68):76–82, 2009.
[51] A. Poggi, D. Lembo, D. Calvanese, G. De Giacomo, M. Lenzerini, and R. Rosati. Linking data to ontologies. *J. Data Semantics*, 10:133–173, 2008.
[52] A. Robinson and A. Voronkov. *Handbook of Automated Reasoning*. The MIT Press, 2001.
[53] B. Rossman. Homomorphism preservation theorems. *J. ACM*, 55(3), 2008.
[54] Y. Sagiv and M. Yannakakis. Equivalences among relational expressions with the union and difference operators. *J. ACM*, 27(4):633–655, 1980.
[55] O. Shmueli. Equivalence of DATALOG queries is undecidable. *J. Log. Program.*, 15(3):231–241, 1993.
[56] M. Y. Vardi. Automata theory for database theoreticans. In *Theoretical Studies in Computer Science*, pages 153–180, 1992.
[57] M. Y. Vardi. Reasoning about the past with two-way automata. In *ICALP*, pages 628–641, 1998.
[58] T. Wilke. Alternating tree automata, parity games, and modal μ-calculus. *Bull. Belg. Math. Soc. Simon Stevin*, 8(2):359–391, 2001.